


NASA/CR-1998-208959



Semiannual Report

April 1, 1998 through September 30, 1998

*Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA*

Operated by Universities Space Research Association



National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

Prepared for Langley Research Center
under Contract NAS1-97046

December 1998

Available from the following:

NASA Center for AeroSpace Information (CASI)
7121 Standard Drive
Hanover, MD 21076-1320
(301) 621-0390

National Technical Information Service (NTIS)
5285 Port Royal Road
Springfield, VA 22161-2171
(703) 487-4650

CONTENTS

	Page
Introduction	ii
Research in Progress	
Applied and Numerical Mathematics	1
Physical Sciences, Fluid Mechanics	19
Computer Science	32
Reports and Abstracts	43
ICASE Interim Reports	52
ICASE Colloquia	53
ICASE Summer Activities	56
Other Activities	61
ICASE Staff	62

INTRODUCTION

The Institute for Computer Applications in Science and Engineering (ICASE)* is operated at the Langley Research Center (LaRC) of NASA by the Universities Space Research Association (USRA) under a contract with the Center. USRA is a nonprofit consortium of major U. S. colleges and universities.

The Institute conducts unclassified basic research in applied mathematics, numerical analysis and algorithm development, fluid mechanics, and computer science in order to extend and improve problem-solving capabilities in science and engineering, particularly in the areas of aeronautics and space research.

ICASE has a small permanent staff. Research is conducted primarily by visiting scientists from universities and industry who have resident appointments for limited periods of time as well as by visiting and resident consultants. Members of NASA's research staff may also be residents at ICASE for limited periods.

The major categories of the current ICASE research program are:

- Applied and numerical mathematics, including multidisciplinary design optimization;
- Theoretical and computational research in fluid mechanics in selected areas of interest to LaRC, such as transition, turbulence, flow control, and acoustics;
- Applied computer science: system software, systems engineering, and parallel algorithms.

ICASE reports are considered to be primarily preprints of manuscripts that have been submitted to appropriate research journals or that are to appear in conference proceedings. A list of these reports for the period April 1, 1998 through September 30, 1998 is given in the Reports and Abstracts section which follows a brief description of the research in progress.

*ICASE is operated at NASA Langley Research Center, Hampton, VA, under the National Aeronautics and Space Administration, NASA Contract No. NAS1-97046. Financial support was provided by NASA Contract Nos. NAS1-97046, NAS1-19480, NAS1-18605, NAS1-18107, NAS1-17070, NAS1-17130, NAS1-15810, NAS1-13394, NAS1-14101, and NAS1-14472.

RESEARCH IN PROGRESS

APPLIED AND NUMERICAL MATHEMATICS

BRIAN G. ALLAN

Reduced Order Model for Sensor/Actuator Problem

Placement of micro-actuators and sensors for active flow control problems is complicated by the distributed nature of the flow equations and the dynamics of the closed-loop system. The goal of this research is to provide insight into the areas of the flow where actuation and sensing are important. This is done by performing an optimal feedback control design assuming distributed actuation and sensing everywhere in the flow field. The resulting optimal feedback kernel will then provide insight into the areas of the flow where actuation and sensing should be performed. Since the calculation of this optimal feedback kernel has a large computational cost, (order N^3 where N is the number of state variables) a reduced-order model must be developed.

The flow is modeled by the unsteady incompressible Navier-Stokes equations in vorticity-stream function form. These equations are then linearized about a desired steady-state flow field and discretized using a finite difference scheme. A reduced-order model is then generated from the linearized flow equations by constructing a Krylov subspace using the Arnoldi process. An implicitly restarted Arnoldi method is used in order to extract the leading eigenvectors from the Krylov subspace while avoiding storage problems. This Krylov subspace produces a reduced-order model of the linearized flow equations used in the optimal feedback control design. We have successfully applied this method on a two-dimensional shear layer and a forced cavity flow problem at low Reynolds numbers. This reduced-order model is then used to calculate an optimal feedback gain using an LQR approach. These feedback kernels then show the spatial regions where sensing and actuation are favorable.

We are currently working on extending this research for the flow around an airfoil.

Active Flow Control for an Airfoil at a High Angle of Attack

Experimental results have shown that the introduction of an oscillating jet at the leading edge of an airfoil can increase its performance at high angles of attack. To improve the performance of this system, and make it more robust to disturbances, the design of a feedback controller will be investigated. In this study, we will develop a model of the airfoil system using a computational fluid dynamics code which will then be used to design a feedback controller.

The approach taken here is to first validate the computational simulations using experimental data. Once the code has been validated, a model of the system will be designed using input/output data generated by the flow solver. This model will then be incorporated into an optimal robust control design. The performance of the controller will be evaluated from closed-loop simulations of the system. Presently, computational simulations of the airfoil with an oscillating jet system have been performed. These computations were done using a compressible Reynolds averaged Navier-Stokes code developed at NASA. These simulations have shown good agreement with experimental data by predicting the performance increase at several mass flow rates.

We plan to continue this work by using the flow solver to develop a model of the airfoil system. Once a model has been developed, a robust feedback controller can be designed and evaluated. This work is done in collaboration with S.S. Ravindran (NASA Langley and NRC).

EYAL ARIAN

Approximations of the Newton Step for Large Scale Optimization Problems

Quasi-Newton methods for large scale optimization problems are powerful but suffer an initial slow convergence rate. Our goal is to develop a new iterative method, for the solution of large scale optimization problems, that will allow a better approximation for the Newton step right from the first optimization steps.

In the course of the optimization process, systems of linear equations are constructed that contain the linearized state operator and its adjoint. These have to be solved at each iteration to achieve convergence of the iterates to the Newton step. We are investigating a defect-correction method to solve these systems of equations for highly ill-conditioned problems with many design variables. Preliminary numerical tests on the potential small disturbance shape optimization problem are promising.

Our plan is to further investigate the above method for applications that are governed by non-linear equations. This approach can be naturally embedded in a SQP formulation of the problem.

This work is done in collaboration with A. Battermann and Prof. E. Sachs (Universität Trier, Germany).

F. SCOTT COWAN

Incorporating Perspectives into the Design of Conceptual Engineering Systems with Living Systems Theory

Complex engineered systems such as aeronautic vehicles consist of many different interfacing subsystems and are subject to numerous constraints and requirements regarding operability, safety, manufacturability, maintainability, economics, etc. The design of such systems requires the teamwork of specialists from a variety of disciplines — each retaining different but necessary perspectives of the system under design. A successfully designed solution depends largely upon the collaboration of the multi-perspective laden team. However, many design teams often experience difficulties stemming from the differences in perspectives. For instance, discipline-based perspectives generally consist of unique knowledge, grammars, vocabularies, and representations; furthermore, human perspectives are subjective and experience-based rather than absolute in terms of correctness, validity, importance, etc. In this work we are developing a new conceptual design methodology for engineering systems, such as aircraft, in which various perspectives are appropriately incorporated. This leads to better creation and integration of higher-fidelity system design models within a concurrent engineering design framework.

To model various perspectives explicitly, we employ the lexicon and representation scheme of Living Systems Theory (LST). LST is the conceptual framework developed to integrate the findings of system theorists in biology, physiology, neurology, the social sciences, economics, and management; it serves as a unified and general systems theory that deals with hierarchical systems. It has also been successfully utilized to model nonliving engineering systems at a functional level of design abstraction. In the context of a Decision-Based Design paradigm, LST is a core lexicon and representation scheme with which different disciplinary and life-cycle perspectives may be modeled, communicated, and integrated to create better understanding among the design team; this should improve the efficiency and effectiveness of designers as they make design decisions (e.g., selections and compromises). A higher-fidelity, holistic systems model

results, appropriate for the synthesis-based activities within conceptual design. We are creating a case study by applying this approach to a NASA Langley project: the design of an airborne differential absorption lidar (DIAL) instrument, of which we have created several perspectual LST models.

Once all perspectual models of the DIAL system are created, we will develop an integration strategy and investigate software implementation. We will also develop a method for refining such qualitative models into more quantitative and concrete design representations, suitable for analysis-based design activities. This work is performed in collaboration with Dr. Russell DeYoung (NASA Langley) and Dr. Farrokh Mistree, Dr. Janet K. Allen, and Dr. David Rosen (Georgia Institute of Technology). Further collaboration and assistance is provided by Ed Dean, Dr. Robert Weston, and Dr. Thomas Zang (NASA Langley).

DAVID L. DARMOFAL

Eigenmode Analysis of Boundary Conditions for the One-dimensional Preconditioned Euler Equations

Local preconditioning has been successfully utilized to accelerate the convergence to a steady-state for Euler and Navier-Stokes simulations. Since preconditioning effectively alters the time-dependent properties of the governing partial differential equation, modifications of the numerical discretization can be required. For example, upwind methods for inviscid problems must be based on the characteristics of the preconditioned equations instead of the unpreconditioned equations. Similarly, the behavior of boundary conditions in conjunction with preconditioning will also be altered. While the effect of preconditioning on boundary conditions is well-known, to date, no quantitative analysis has been performed. The purpose of this work is to analyze the effect of preconditioning on several different boundary conditions commonly used in numerical simulations.

We consider the one-dimensional, preconditioned Euler equations linearized about a steady, uniform, subsonic mean state and calculate the exponential decay rates for perturbations under different sets of boundary conditions. The work is an extension of the analysis of Giles for the one-dimensional, unpreconditioned Euler equations. Boundary conditions based on the Riemann invariants of the unpreconditioned Euler equations are found to be reflective in conjunction with preconditioning. The problem is most detrimental at low Mach numbers where the decay rate of perturbations approaches zero. Boundary conditions which specify entropy and stagnation enthalpy at an inflow and pressure at an outflow are found to be non-reflective with preconditioning. Numerical results were presented which are in good agreement with the analytic predictions.

An ICASE report is in preparation documenting the analysis and numerical results.

BORIS DISKIN

Analysis of Convergence of Defect Correction Iterations for Convection Equations

Convergence properties of the defect correction method applied to elliptic equations are now fairly well understood and quantitative predictions of efficiency of a given scheme agree very closely with numerical calculations. The behavior of the method in application to hyperbolic equations is much less clear. In this period, we systematically studied the convergence of the defect correction iterations for a model convection equation corresponding to flow at some angle of attack to a uniform Cartesian grid. This research was motivated by the search for an explanation of convergence properties of a full Euler system solver.

A comprehensive half-space mode analysis closely predicting convergence properties of the model problem has been developed. On the base of this analysis, we have explained many surprising details observed

in numerical calculations (e.g., establishment of a good asymptotic convergence rate after many poorly converging iterations). A new measure for the discretization accuracy of a given discrete operator (called the penetration distance) connecting the operator approximation order, the frequency of the incoming oscillations and the angle of attack has been introduced. It has been found, analytically and experimentally, that the convergence properties of the defect correction iterations are determined by the difference in the penetration distances of the discrete operators involved. The obtained analytical results provide a quantitative prediction for the number of iterations required to converge solution to within discretization error. They also predict the residual convergence rate on different stages of the defect correction iterations. All the predictions are in very close agreement with numerical results. This analysis has also initiated several recent proposals aiming to improve the convergence properties (e.g. use of a second order operator to solve the correction equation, involving a predictor-corrector method and accelerating the convergence by a separate multigrid method).

We are going to extend this analysis to the full system of Euler equations. We also plan to implement some of the proposed ideas in the framework of the existing 3D flow solver.

This work was done in collaboration with J.L. Thomas (NASA Langley).

JACK R. EDWARDS

Investigation of Optimally Convergent, Preconditioned Multigrid Solvers

The search for optimal convergence for Navier-Stokes calculations is turning toward a re-examination of implicit schemes based on approximate factorization, alternating-direction implicit, or relaxation techniques. Newer procedures utilize subiteration techniques to reduce factorization error and are embedded within an overall multigrid framework, where they function as a powerful smoother. Even with these procedures, convergence for low-speed flows (free-stream Mach numbers less than about 0.15) may be highly non-optimal due to wide variations in characteristic wave speeds. One means of alleviating this is through the concept of time-derivative preconditioning, which rescales acoustic waves such that the condition number remains bounded as the Mach number approaches zero. As preconditioning changes the eigensystem of the Euler equations, every facet of a CFD code – discretization, integration, and boundary condition implementation – is affected by its inclusion. The primary purpose of the author's stay at ICASE was to implement a time-derivative preconditioning strategy into CFL3D, a multi-purpose Navier-Stokes solver, and into CFL3D-DDADI, a highly implicit variant of CFL3D under development by J. Thomas and co-workers in the Aerodynamic and Acoustic Methods branch at NASA Langley.

The objective of this study has been met. The Weiss-Smith preconditioning matrix (a structurally-simple variant of preconditioners proposed by Turkel, Choi, and Merkle) has been implemented into CFL3D and CFL3D-DDADI. A Roe-type upwinding method based on the preconditioned system has also been derived and implemented, along with an approximate linearization of the Roe method. CFL3D calculations for Euler flow over an airfoil and Navier-Stokes flow over a flat plate have demonstrated Mach-number independent convergence rates. CFL3D-DDADI calculations of Euler flow over a bump in a channel have demonstrated Mach-number and grid-independent convergence rates. Typical multigrid convergence factors are between 0.1 and 0.5, indicating that the CFL3D-DDADI scheme, enhanced by time-derivative preconditioning, is nearly optimal in terms of performance.

Work to be completed includes the development of diagonalized ADI techniques based on the preconditioned system (the preceding results were obtained for non-diagonalized ADI schemes), the testing of the CFL3D-DDADI method for more complicated flows, and the investigation of methods for extending the CFL3D-DDADI procedures to three dimensions without sacrificing performance.

This work has benefited from helpful discussions with Jim Thomas, Eli Turkel, Rolf Radespiel, and Bram van Leer.

MICHAEL LEWIS

Pattern Search Methods for Nonlinear Optimization

Pattern search methods for nonlinear optimization have a number of features that make them attractive for use in engineering optimization. These methods are easy to understand and implement, are scalably parallel, and neither require nor estimate derivatives.

In recent work, we have developed the first pattern search algorithms for general nonlinearly constrained optimization guaranteed to possess first-order stationary point convergence. We have now begun an implementation of the new, general classes of pattern search algorithms. This new implementation will allow us to investigate various algorithmic approaches, as well as opportunities for improved computational parallelism.

Among the algorithmic approaches we will investigate are techniques to improve scaling in pattern search algorithms via the aggregation of similarly scaled design variables. This work on pattern search is being done in collaboration with Elizabeth Dolan, Michael Trosset, and Virginia Torczon (The College of William & Mary), and William Hart (Sandia National Laboratory).

A Posteriori Finite Element Bounds for Sensitivity Calculations

In the optimization of systems governed by differential equations one would like to use the coarsest mesh possible at any given step so as to reduce the cost of the optimization iteration. In a recent series of papers, Patera, Peraire, and their collaborators have presented an *a posteriori* approach to computing *quantitative* bounds on the mesh dependence of certain functionals of the solutions of differential equations. We have begun to apply these ideas in the context of optimization.

We have developed *a posteriori* bounds for sensitivities of output linear functionals with respect to various parameters (such as coefficients) in boundary-value problems. Using either the sensitivity equations or adjoint equations one can write the output's sensitivity as a functional of the solution of a system of differential equations. One then computes bounds on the error in the sensitivities on a coarse grid relative to a finer grid. Numerical results indicate that the bounds can be quite good. We have also extended the *a posteriori* bound approach to certain non-smooth functionals.

We are currently implementing an approach to using the *a posteriori* bound procedure in connection with pattern search methods, a first step in a larger investigation of using approximate function values with error bounds in optimization.

This work was done in collaboration with Tony Patera and Jaime Peraire (Massachusetts Institute of Technology).

IGNACIO M. LLORENTE

Plane-Smothers for Multi-Block Grids

Standard multigrid methods are not well suited for problems with anisotropic coefficients which can occur, for example, on grids that are stretched in order to resolve a boundary layer. There are several different modifications to the standard algorithm that yield efficient methods for anisotropic problems. One of the more efficient approaches is the application of robust multigrid smoothing processes with standard

coarsening (use of alternating-line relaxation in 2D and alternating-plane relaxation in 3D). However, this may be difficult to implement in codes using multi-block structured grids because there may be no natural definition of global line or plane. Block structured grids are commonly used in NASA fluid dynamics codes to deal with complex geometries and facilitate parallel processing. This inherently limits the range of an implicit smoother to only the portion of the computational domain in the current block. The goal of the present research is to study in detail, both numerically and analytically, the behavior of plane implicit smoothers in multi-block grids in order to provide guidance to practitioners using these block-structured grids.

We have developed a flexible 3D code to study the behavior of blocked plane smoothers. The current version of the code solves the nonlinear diffusion-convection equation by using a full multigrid approach with the full approximation scheme for V-cycles. The equation can be solved in a multi-block grid built by joining cell-centered blocks with grid stretching to solve complex geometries. The blocks can overlap with neighboring blocks in order to reduce the negative effect of the blocking on the convergence rate. The code is implemented in Fortran77 and has been parallelized using the Silicon Graphics, Inc. `doacross` directives for shared memory parallel computing. Results are encouraging: textbook multigrid convergence rates can be obtained with an overlap just 10% of the number of cells in the block and the convergence rate can be maintained by increasing linearly the number of overlapping cells with the problem size. The multigrid problem tends to be a domain decomposition problem for high anisotropies. The results obtained thus far show alternating-plane smoothers to be very robust even in multi-block grids.

A report will be written with the numerical results and an analytical study to verify these results. The code updates the blocks following a Gauss-Seidel ordering. However, in order to run the code on a parallel computer, so that each processor solve a set of blocks, the update ordering must present a higher parallelism grade. The next step will be to check the convergence rate of red-black Gauss-Seidel ordering of the blocks. We plan to continue work on block-structured algorithms. In particular, we will study the applicability of blocked alternating-plane methods as multigrid smoothers for convection dominated problems and more complicated PDE's.

This work has been done in collaboration with N.D. Melson (NASA LaRC).

JOSIP LONČARIĆ

Spatial Structure of Optimal Flow Control Around Airfoil

Designing distributed control systems begins with the sensor/actuator placement problem. While in some situations discrete search of combinatorial complexity seems unavoidable, continuum problems suggest solving a related question. *If one could sense everything and actuate everywhere, what should one do?* The answer to this question has polynomial complexity (of order N^3 where N is the number of state variables) and can serve as the initial effectiveness filter capable of rejecting a large portion of the design search space. This favorable situation can have several causes. In our investigation we focus on the effect of no-slip boundary conditions on an optimal flow control of the unsteady Stokes flow around the NACA 0015 airfoil. This test problem aims at the development of numerical methods capable of solving the problem of stabilizing the desired flow around wings at low to moderate Reynolds numbers.

Our approach begins by defining a *pseudodifferential representation* ξ of the flow. This step is similar to the Wiener-Hopf factorization of the Laplacian. We have shown that for the flow around a cylinder this leads to an explicit diagonalization of the system dynamics $\xi_t = \mathcal{A}\xi$ by means of the Fourier and Weber transforms. We then pose and solve an optimal distributed LQR problem with gain ϵ . A rational approximation to the

optimal feedback kernel is derived and shown to perform within 0.026 percent of the exact optimum even in the worst case. Using the vorticity representation in conformally mapped geometries, this approximation is decomposed into the analytic free space solution and a boundary term which can be evaluated numerically.

A numerical procedure to display the spatial form of the nearly optimal feedback in response to a disturbance vortex has been developed. While the local vorticity damping is a part of this solution, the boundary term part of the optimal actuation is of particular interest. The insight gained in this study will provide guidance for sensor/actuator placement as we extend the control theoretic approach to low and moderate Reynolds number flows.

RICHARD LONGMAN

Improved Speech Coding Based on Open-loop Parameter Estimation

Speech coding converts speech into a coded form which can be transmitted. Then the receiver uses the transmitted information to reconstruct the speech. The effectiveness of a speech coding algorithm is a function of how many bits must be transmitted, and the quality of the reconstructed speech. Linear speech coding develops a linear difference equation model of successive segments of the speech signal and transmits the coefficients of the model for each segment, and also the equation error residual for each time step. A reduced number of bits is used in the residual, and this accomplishes a compression of the number of bits that need to be transmitted. In previous work done at ICASE, a linear predictive speech coding algorithm was developed that not only optimizes the linear model coefficients for the open loop predictor, but does the optimization including the effects of limiting the bits in the transmitted residual. It also simultaneously optimizes the quantization levels used for each speech segment. The purpose of the research reported here is to improve this algorithm, producing substantially improved quality of the reconstructed speech.

The optimization involved in the algorithm is a nonlinear minimization problem, and can converge to local minima. Here a method is developed based on correlation functions of the data, in order to produce good initial starting data. In addition, the number of bits in the transmitted residual is monotonically increased in a procedure analogous to continuation methods. Examples of speech encoding and decoding are given for eight speech segments, and signal to noise levels as high as 47 dB are produced, which corresponds to very high quality speech. This is a dramatic improvement over the performance of the original version of the algorithm.

As is standard in linear speech coding or LPC, the optimization is done on the open loop speech analysis model rather than directly optimizing the quality of the reconstructed speech, because this optimization is much simpler. Here we also study the potential for improvement in speech quality by doing this direct optimization. Typical nonlinear optimization codes for nonlinear least squares problems are prohibitively slow for this investigation. A specialized code was developed using knowledge of both the gradient and the Hessian in order to make this investigation. The results suggested that the speech coding algorithm developed here optimizing the open loop speech signal is nearly optimal for the final reconstructed speech signal as well. Thus, the examples suggest that the algorithm here is close to giving the best performance obtainable from a linear model, for the chosen order with the chosen number of bits in the codebook.

A draft document reporting these results has been produced and is under review as a NASA TM. This work was done in collaboration with Dr. Jer-Nan Juang (NASA Langley) and Ms. Ya-Chin Chen (ICASE).

Improved Linear System Identification

The OKID algorithm for developing state space models from input-output data was developed at NASA Langley by Drs. Jer-Nan Juang, Minh Q. Phan, Lucas Horta, and Richard Longman. The algorithm was demonstrated to be very effective on difficult high dimensional spacecraft structural dynamics problems. A program package is distributed by NASA, and the method is used by over 100 industries and universities. The algorithm first produces an observer, which under appropriate conditions is a Kalman filter, and from the observer it then obtains the system model. Equation error is minimized in producing the observer. Under the conditions associated with a Kalman filter, i.e., white additive plant and measurement noise, the residual becomes white and the bias in the model identification disappears. This is an important advantage of OKID. However, in real data on spacecraft structures where the true system is infinite dimensional, experience has shown that the damping of the identified structural modes is consistently overestimated. The objective of the current research is to develop an algorithm that minimizes the output error instead of the equation error, and the numerical results show that the phenomenon of bias in the identified damping is eliminated.

The problem of minimizing output error is a nonlinear optimization. The OKID result is used for starting values in a sequential quadratic program (SQP), so that the program iterates starting from the minimum equation error and converges to minimum output error. Several linear model types are considered, including a general state space representation, autoregressive models with an exogenous input (ARX), and second-order modal models appropriate for spacecraft and aircraft structural modeling. These problems can easily be high dimensional, and one must use as much information as possible to achieve fast convergence rates. For each type of model, analytical methods are developed that determine both the gradient and the Hessian. The code developed is optimized for efficiency, making use of many interesting analytical relationships to speed up the computation. A nonlinear least squares algorithm is generated that evaluates the eigenvalues of the second derivative portion of the Hessian and replaces any negative eigenvalues by zero. The running time for the code is quite reasonable, whereas without the intricate development of the analytic derivatives, the computation time could easily be prohibitive. The resulting code is of immediate usefulness to those involved in modal testing, producing improved structural dynamics models. And it has quite general usefulness in producing improved models in any linear system identification problem.

This work was done in collaboration with Dr. Jer-Nan Juang; (NASA Langley). Written documentation of the results is in preparation and will appear as a NASA TM.

Approaches to General Purpose Input-output Modeling from Non-linear System Data

For linear systems, there are general purpose formulations such as modern state space representations that can be used for all linear systems. General purpose codes such as OKID can take rich input-output data for a linear system and develop a mathematical model that reproduces the system input-output behavior. Such models are then used for control system design. For nonlinear systems it is not straightforward to find a general purpose representation. And, in many situations, a representation based on physical understanding to supply the nonlinear terms can require an unreasonable amount of time to develop and an unreasonable amount of time to use. In the problem of flow control past air foils, physically based models result in large finite element codes. For purposes of control system design, such models are far too complex for direct use. If simpler models could be developed to reproduce the observed relationship between actuator motions and measured responses, the control law design process could be greatly facilitated. The objective of this research is to investigate the types of general purpose models available, develop understanding of their relationships, and propose approaches to development of general purpose nonlinear system identification algorithms.

One class of nonlinear models is Volterra series. Another class is nonlinear ARX models. We show how these are related, and how one can be converted to the other. This gives insight into how one should approach the identification process to minimize the number of nonlinear terms that are needed to model the nonlinear behavior. A procedure is developed to first identify the linear part of the behavior, and isolate the part of the input-output data that must be captured by nonlinear effects. A method is suggested to rank the effectiveness of candidate nonlinear terms in a nonlinear ARX, or a still more general model allowing a choice of still more general nonlinear terms. The proposed procedure is expected to facilitate the development of nonlinear models for complex systems, doing so directly from data without requiring physical modeling. And the model developed should be far simpler, requiring far fewer terms, than the Volterra series approach.

A working document outlining the proposed approach to general purpose nonlinear system identification has been written. It is intended that code will be developed at NASA and then used to investigate the effectiveness of the approach on difficult identification problems. This work was done in collaboration with Dr. Jer-Nan Juang (NASA Langley).

VALERIE M. MANNING

Collaborative Optimization Implementation of Aerospoke Nozzle Design

A model of an aerospoke rocket nozzle that consists of coupled aerodynamics, structural and thermal analyses has been developed. The model will be used to demonstrate multidisciplinary design optimization (MDO) capabilities and assess the performance of various MDO approaches. In this work the aerospoke nozzle problem is implemented in the collaborative optimization framework.

Before attempting to perform MDO on any design, the problem must be thoroughly decomposed. The aerospoke nozzle is a highly coupled system in which aerodynamics, structural, and thermal effects all contribute to the design. Collaborative optimization (CO) is a method which serves to eliminate direct interaction between disciplines, and rather channels all information flow to and from a single, governing "system level" optimizer. The case of the aerospoke nozzle is a complicated one to decouple in such a manner, each discipline strongly influences the mechanical sizing and loading and the thermal sizing and loading. CO allows each discipline to perform its local optimization in parallel with the others, without direct concern for what is going on elsewhere. This is due to the nature of engineering design in industry, where human designers are grouped by specialty, often in physically distinct locations. In order to determine exactly what the interactions are for this problem, personnel from NASA Langley Research Center (Structures), NASA Ames Research Center (Aerodynamics), and Rocketdyne Division of Boeing North American, Inc. (Thermal) were sequestered to settle upon the required interactions. As many of the shared variables are arrays (pressure distributions, nozzle deflections), reduced basis modeling (RBM) approaches have been proposed, and the suitability of various fit types were discussed. Presently, the problem has been decomposed and defined in a CO framework, and discussions with designers in each discipline to ensure that the required information is being passed and that RBM assumptions are sound have been completed.

Upon the decomposition of this aerospoke nozzle design problem into other viable MDO frameworks has been completed, a method will be chosen to implement and perform the optimization.

This work was done in collaboration with Tom Zang (NASA Langley).

Multidisciplinary Design and Optimization of a Natural Laminar Flow SST

The overall objective of the work is to develop the design process of a next generation high speed aircraft, specifically a supersonic transport created to naturally take advantage of laminar flow. There are many

pieces to the puzzle. Previously, a multidisciplinary design optimization (MDO) integrating aerodynamics, structures, and mission was performed on an SST wing. The current goal is to extend this to a wing-body concept in order to capture the effect of the fuselage on the wing pressure distribution. Therefore, a method to allow body design in addition to wing design requires automatic grid generation of a fuselage for use with a surface panel method as a function of a given set of design variables.

Many grid generation options exist, from CAD programs to codes created specifically for use with CFD analysis. The difference in the present case is that the entire grid must be generated from a small set of design variables at each iteration of an optimization process, and a FORTRAN source is desired. Optimization means that design flexibility and speed are each very important. Given, for example, a fuselage length, nose and tail fineness ratio, fuselage upsweep (for take-off clearance) and the amount of supersonic area-ruling desired, a grid must be created and intersected with the wing model in such a way to satisfy the sensitive 3D aerodynamics and FEM codes. Use of an aerodynamics code means generation of wakes from the wing and body as well. This process was automated, in part with the help of BINTER, a geometry package for generation of input data for 3D potential flow programs (NASA CR 2962). Once a coarse gridding from the design variables is performed, BINTER does finer re-gridding, and also computes intersection curves of the wing and fuselage, returning the new, intersected grid. Once this process was fully automated, it was embedded into the aerodynamics optimization package.

The next step in the aerodynamics portion of this work is to incorporate a viscous analysis package which consists of a stripwise 2D boundary layer code and an empirically based crossflow Reynolds number computation function (previous work). Once this is completed, along with enhancements in the disciplines of aerodynamics and structures, the full optimization will be performed.

This grid generation package could not have been developed without the help of Bob Weston (NASA Langley).

DIMITRI J. MAVRIPLIS

Parallel Implementation of a Three-dimensional Unstructured Multigrid Solver

Unstructured mesh Navier-Stokes solvers offer great potential for reducing the turnaround time associated with complex geometry aerodynamic analysis. However, the large computational overheads associated with unstructured mesh methods requires the use of efficient solution algorithms which can be ported to massively parallel architectures. The purpose of this work is to implement a previously developed anisotropic unstructured multigrid solver to massively parallel computer architectures, and to benchmark the performance of this implementation on large scale analysis runs.

The fine and coarse levels of the unstructured multigrid algorithm are all partitioned sequentially before being distributed on the target parallel machines. Because the algorithm makes use of implicit line solves, the partitioning must be executed in such a way that the implicit lines of the various mesh levels are not intersected by processor boundaries. This is achieved by contracting the mesh graph along the implicit lines and partitioning the contracted (weighted) graph rather than the original graph of the mesh, which is then used to infer the final mesh partition upon de-contraction. The communication patterns (which remain static for the duration of the analysis) are then precomputed and stored. The implementation of the parallel solver is based on the MPI communication primitives. The computer platforms consisted of a 512 processor Cray T3E, and a 128 processor SGI Origin 2000. Access to both hardware platforms in dedicated mode was provided by Cray Research, Eagan, NM.

Good scalability of the unstructured mesh multigrid solver has been demonstrated on both hardware platforms for problems involving several hundred thousand grid points, up to 13.5 million grid points, with near perfect scalability being achieved for problems as small as 2 million points. A high-lift component geometry case has been solved on a grid of 13.5 million grid points in 4 hours on 512 processors of the Cray T3E. This represents the largest unstructured Navier-Stokes case attempted to date.

Future work is concentrated on improving the numerical efficiency of the unstructured multigrid algorithm, as well as demonstrating a full aircraft high-lift configuration analysis capability on these machines using upwards of 20 million grid points.

ROLF RADESPIEL

Analysis and Refinement of Recent High-resolution Schemes for Transonic and High-speed Flow

Recently developed dissipation models based upon flux-vector splitting promise perfect resolution of grid-aligned contact discontinuities and shock capturing with a single interior point. Jameson's HCUSP and Liou's AUSM are representatives of these schemes. In the present work we analyze and improve the HCUSP dissipation scheme because it combines easily with multistage time stepping that we use for smoothing in the multigrid context. The performance of HCUSP is assessed in comparison with matrix dissipation schemes that are well-known for their accuracy and efficiency. The present work shall yield a statement whether HCUSP is to replace existing matrix dissipation schemes for reasons of accuracy, robustness and convergence.

The basic HCUSP scheme has been improved in two ways: 1) The underlying first-order scheme is extended to treat shock waves oblique to the grid without oscillations and 2) the higher-order extrapolation is extended to avoid undesired limiting in smooth regions of the flow and to reduce overall level of background dissipation. Our 2D viscous results indicate comparable grid and residual convergence for HCUSP and matrix dissipation. 3D transonic comparisons are underway. HCUSP, however, is more robust in capturing strong discontinuities that occur in hypersonics. Here, we have successfully completed 2D and 3D computations. Applications of HCUSP to the 2D and 3D Navier-Stokes equations show a significant slowdown of the convergence rates when compared to the Euler results (similarly as for matrix dissipation). Hence we have devised preconditioners for the residual to be used in connection with multistage schemes. First analysis reveals that our pointwise preconditioners will probably not improve the convergence rate for viscous flows. Hence there remain differential preconditioners in order to remove stiffness due to cell aspect ratio. We have completed the coding of the HCUSP preconditioners in 2D and have begun the debugging phase. The work performed thus far indicates strong potential of the flux-vector splitting scheme HCUSP with respect to accuracy and robustness whereas the potential for rapid convergence remains to be seen.

In the future we plan to investigate the performance of differential preconditioners in order to remove stiffness associated with Navier-Stokes computations. The work is done in collaboration with E. Turkel (Tel-Aviv University), R.C. Swanson (NASA Langley), and V. Vatsa (NASA Langley).

C.-C. ROSSOW

A Flux Splitting Scheme for Compressible and Incompressible Flows

The requirements on discretization methods suitable for the simulation of technically relevant flows cover a very broad spectrum, and are sometimes even contradicting: in order to obtain accurate solutions, only a

minimum amount of artificial viscosity can be tolerated, especially for the resolution of boundary layers. For flows with strong shock waves, however, robustness becomes of primary importance, especially near vacuum conditions where the prediction of positive quantities like pressure and density is likely to occur. In such situations, the proper amount of numerical viscosity needs to be supplied to achieve a converged solution. In order to establish adequate discretization schemes, recently the development of numerical schemes focuses on the construction of hybrid flux-splitting formulations, where the robustness of flux-vector-splitting is to be combined with the accuracy of flux-difference-splitting. Based on considerations for schemes like AUSM, CUSP, and LDFSS, the author derived a simple flux splitting method (MAPS) for compressible flows, which is both accurate and robust, despite its simplicity. Besides the derivation of numerical methods for the computation of compressible flows, over the past several years considerable effort has been spent on the extension of schemes designed for compressible flow to solve incompressible flow problems. This effort is driven by various reasons. First, flow problems may be composed of compressible and incompressible regions, such as the flow around high lift devices. Another example is flow with surface heat transfer or volumetric heat addition, a flow which is compressible despite low speeds. A further reason is the necessity to establish numerical methods handling the broadest range of flow conditions in order to avoid dealing with multiple flow codes.

In order to establish a method suitable for both compressible and incompressible flows, the well-known flux-difference-splitting scheme was expanded in functions of Mach number. This expansion yielded both physical significance of the dissipative terms and a subdivision into terms either important for compressible or for incompressible flow. It was found that the relevant disturbances are given by pressure, normal velocity, and advected species. An inspection of the disturbances active in the compressible flow regime revealed that these are the terms which are incorporated into the MAPS scheme. Therefore, in order to extend MAPS to incompressible flows, the incompressible terms identified in the flux-differencing-scheme were incorporated into MAPS. The incompressible terms mainly consisted of pressure differences, thus providing a pressure coupling when the Mach number approaches zero. Using the preconditioner of Choi and Merkle, the matrix only slightly modified to ensure a consistent transition into the compressible flow regime, the resulting method proved to be capable of computing airfoil flows with convergence rates independent of the freestream Mach number. It was observed that almost no adaption of numerical parameters was necessary when computing high lift flows near separation at incompressible conditions. Furthermore, the resulting scheme was capable of resolving hypersonic flows without encountering numerical problems as known from flux-difference-splitting.

The scheme developed is a typical upwind-scheme which relies on the construction of variables at cell interfaces. It is worthwhile to consider a formulation with second and fourth differences analogously to the Matrix-dissipation. Due to the close relationship of the flux-differencing approach, which is the basis of the matrix dissipation, the adaption of the MAPS scheme to such a formulation should be feasible. Further research needs to be directed towards suitable limiter functions in case the upwind-formulation needs to be retained, in order to avoid convergence problems. An interesting field of further research will be the incorporation of the MAPS discretization into the framework of a method based on a Least Squares approach.

DAVID SIDILKOVER

Factorizable Schemes and Essentially Optimal Multigrid Solvers for the Flow Equations

The main objective of this work is to develop discretization schemes that facilitate construction of the essentially optimal multigrid solvers for the equations of steady compressible flow. Our first targets are the

Euler equations in two dimensions. However, the methodology being developed is very general. It can be extended to Navier-Stokes equations and to three-dimensional problems.

A genuinely multidimensional high-resolution discretization for the steady compressible Euler equations was constructed recently. One of its advantages is that it allows to combine an h -ellipticity and high-resolution properties in the same scheme. Another fundamental advantage is that this approach also leads to schemes that are *factorizable*, i.e., allow to distinguish between the advection and full-potential factors of the equations. This property is crucial for constructing essentially optimal multigrid solvers. The constructed discretization approximates equations in their usual conservative form.

The current work is devoted to the construction of the multigrid solver, that utilizes the factorizability property of the scheme. Extensions of the approach to viscous and three-dimensional problems are in work as well.

Towards Genuinely Multidimensional Schemes/Solvers for Unsteady Problems

The previously developed genuinely multidimensional approach towards discretizing the flow equations addressed the steady case: the resulting schemes are second-order accurate at the steady-state only. Extension of this approach to unsteady problems is not straightforward. The main motivation for doing this is expectation that it will result in discrete schemes that facilitate the design of very efficient implicit solvers. This is due to the h -ellipticity and factorizability properties of the discretizations.

It was necessary to revisit the one-dimensional case first and to design a new one-dimensional high-resolution scheme that is compatible with the genuinely multidimensional approach. The resulting scheme becomes implicit if the Courant number is larger than one and is quasi-explicit (i.e., can be dealt with as explicit) otherwise. At the first stage, a quasi-explicit case was studied in details. It has been proven that the scheme satisfies the TVD property. An extensive set of numerical experiments was conducted to verify the accuracy and good discontinuity capturing capabilities of the discretization. A two-dimensional genuinely multidimensional time-space accurate scalar advection scheme was constructed as well.

The next natural step will be to consider the implicit 1D case and to construct a fast implicit solver. Then we plan to address the 2D unsteady case. This study is expected to result in very efficient solvers for unsteady problems. In particular this approach should provide some very interesting possibilities for treatment of the low speed flow.

BAMBANG I. SOEMARWOTO

Multi-disciplinary Analysis through Optimization

Mono-disciplinary analysis tools have matured to the stage that allows routine industrial use. These tools, each of which is associated with a different engineering discipline, have resulted from separate developments in the past. The consequence is that, to solve a multi-disciplinary problem, an iterative procedure must be devised such that the mono-disciplinary tools can interact with each other. A Gauss-Seidel type of procedure, or a fixed-point iterative method, is commonly used to obtain the multi-disciplinary solution. However, this may require a severe under-relaxation which leads to many mono-disciplinary analyses. An alternate method is investigated in this research in order to lower the number of mono-disciplinary analyses. This method employs an optimization strategy, in which the minimum of the objective represents a condition for the multi-disciplinary solution.

In this research, the model problem is defined by a rigid weightless airfoil in two-dimensional viscous transonic flow, with the leading edge attached to the plane of the flow by a torsional spring. Given an

angle of attack α , the aerodynamic forces give rise to the pitching moment around the leading edge. In turn, the torsional spring responds with an angular deflection. In the fixed-point iterative method, the angular deflection is fed back as α to the flow solver, and the above process is repeated. In the method investigated in this research, the angular deflection is identified as α^* . By minimizing the deviation between the aerodynamic input α and the structural response α^* , the multi-disciplinary solution has successfully been obtained with a significantly lower computational effort than that of the fixed-point iterative method.

The future plan is to investigate the applicability of the method to the aeroelastic problem of a wing, where bending and torsion as the structural responses alter the shape of the wing surface. An important issue is the geometrical compatibility between the aerodynamic shape of the wing and the structural response. After this has been resolved, the methodology will be embedded in a procedure for multi-disciplinary sensitivity analysis and optimization.

SHLOMO TA'ASAN

Derivation of Macroscopic Equations from Molecular Dynamics Simulations

It has been well known for a long time that the Navier-Stokes equations do not describe well the dynamics of gas at very low density. The common practice is to use direct simulation Monte-Carlo (DSMC) methods in these regimes. It is also well known that the failure of the NS equation in these regimes is not due to the failure of the continuum, but due to inability to express stresses and heat fluxes in terms of low order terms. A similar situation exists in many other fields of science and engineering where different levels of modeling are used to describe relevant physical phenomena. A detailed microscopic model is used to study small scale behavior while a continuum approach is used to model the behavior on large scales. The relation between those different levels of representation is usually not clear and sometimes the continuum model does not correspond exactly to its atomistic/molecular counterpart. From a theoretical viewpoint, the relation between models on different scales is a very difficult problem, and only for a few simple models it is well understood. Computation can aid at this point by deriving the macroscopic equations from the more fundamental level of description, i.e., microscopic levels.

We began our study with Brownian motion where the theory is well understood and develop numerical schemes for the derivation of macroscopic equations from microscopic levels. We have applied the idea to study self-diffusion in gases and were able to recover well known results from statistical physics (related to Einstein's formula for self-diffusion). We then moved to study hydrodynamics effects in gases and focused on the rarefied regime. As expected, discrepancies between molecular level simulation and the NS have been observed. We have engaged in a detailed study of the relation between certain gas models, i.e., hard spheres and Lennard-Jones gases, and the NS equations at different Knudsen numbers. Certain correction terms have been identified, however more study is needed to get a continuum model that will extend the NS equations to rarefied regimes.

We plan to continue this study with the goal of discovering correction terms to the NS equation that will make them valid in a wider range of Knudsen numbers. We anticipate that the resulting macroscopic equation can replace DSMC in certain applications, thus achieving a significant computation speedup.

CHRIS TOOMER

Sensitivity-based Approximation Methods for Use in Multidisciplinary Design Optimization Studies

Work has continued on the development of a direct sensitivity code which updates the flowfield as it marches through a design space. This type of code is quicker than a CFD code but less accurate in the shock area of the flow. However the aerodynamic coefficients are accurate enough to be used in optimization studies, and the flowfields provide good qualitative and in most cases, quantitative data.

During this research period, two approaches for evaluating the viscous flux gradients have been coded and tested. One algorithm produces significantly more accurate and faster results than the other, and has thus become the default approach in the code. The viscous code has been fully tested, and although further improvements will be made to the data storage within the code, the code is ready for release to the users. The inviscid flow sensitivity code already had a 2D aerofoil mesh generator using B-splines for which excellent results had been obtained. The design parameters were the position of the B-spline vertices defining the surface. In the last three months a second 2D mesh generator has been incorporated into the code allowing the user to vary aerodynamic parameters such as camber and thickness. A 3D multiblock structured grid generator for wings has also been successfully incorporated into the sensitivity code. This enables sensitivity studies to be performed changing the wing parameters (e.g., camber, thickness, and twist). Testing has been completed and the code will be available to the users in November.

The future work involves three projects. Firstly, the viscous and grid sensitivity versions of the code will be combined. Secondly, the algorithm will be improved to cope well with fast moving shocks. There are cases where the shock moves 30-40% of the length of the aerofoil section in about one degree change in the angle of attack. The sensitivity code produces good results with such cases in 3D but not so well in 2D, where the shock position can be out by one or more cells. Finally, further work will be undertaken optimizing wing and wing-body configurations. The aim is to investigate which response surface methods should be incorporated into the CFD and sensitivity suite of codes for aerodynamic shape and operating condition optimization.

SEMYON TSYNKOV

Artificial Boundary Conditions for Aerodynamic and Aeroacoustic Computations

Many typical problems in aerodynamics including those that present immediate practical interest, e.g., flows around aircraft, are formulated on infinite domains. It is, however, obvious, that any discretization used for solving such problems on the computer must be finite. Therefore, any numerical solution methodology for these problems has to be supplemented (or, rather, preceded) by a special technique that helps create such finite discretizations. A widely used approach to this problem is based on truncating the original flow domain prior to the actual discretization and numerical solution. Subsequently, one can construct a finite discretization on the new bounded computational domain using one of the standard techniques: finite differences, finite elements, or other. However, both the continuous problem on the truncated domain and its discrete counterpart will be subdefinite unless supplemented by the appropriate closing procedure at the external computational boundary. This is done by using artificial boundary conditions (ABC's); the word "artificial" emphasizing here that these boundary conditions are necessitated by numerics and do not come from the original physical formulation.

As a part of the ongoing research effort on constructing advanced external boundary conditions, we have developed and implemented new global ABC's for computation of flows with propulsive jets. The

algorithm is based on application of the difference potentials method (DPM). Previously, similar boundary conditions have been constructed for calculation of external compressible viscous flows around finite bodies. The foregoing modification for jets substantially extends the applicability range of the DPM-based algorithm. The particular configuration that we have analyzed using new ABC's is a slender three-dimensional body with boat-tail geometry and supersonic jet exhaust in a subsonic external flow under zero angle of attack. Similarly to the results obtained earlier for the flows around airfoils and wings, current results for the jet flow case corroborate the superiority of the DPM-based ABC's over standard local methodologies from the standpoints of accuracy, overall numerical performance, and robustness. The findings of this study are going to be submitted to the 14th AIAA CFD Conference.

Future research in the framework of this project will primarily concentrate on the development of ABC's for the time-dependent problems of wave propagation type. The project is a collaborative effort with Veer N. Vatsa (NASA Langley) and is sponsored by the Director's Discretionary Fund. This work was done in collaboration with Saul Abarbanel (Tel-Aviv University, Israel), Jan Nordstrom (The Aeronautical Research Institute of Sweden), and Viktor Ryaben'kii (Russian Academy of Sciences).

BRAM VAN LEER

Re-dependent Multistage Schemes for Navier-Stokes

Multigrid-type convergence, i.e., convergence in $O(N)$ operations for a solution with N unknowns, can only be achieved if all ingredients of the convergence-accelerating algorithm are working together. It has been demonstrated for explicit Euler discretizations that each of the following ingredients needs to be included:

- local preconditioning;
- multistage method optimized for high-frequency damping;
- semi-coarsened multigrid relaxation.

Progress in converging Navier-Stokes solutions is slow, because the analysis is so much more complex, even with the use of symbolic manipulation. Local preconditioners that are optimal for all Mach and Reynolds numbers have not yet been found, and multi-stage methods await optimization with respect to the Reynolds number.

Using the theorem that the contour $|P(z)| = 1$ of a Chebyshev polynomial encloses an elliptical region, we have been able to construct multistage methods optimized for the *Re*-dependent shape of the Fourier footprint of discrete Navier-Stokes operators. The idea is to fit the footprint in an ellipse, then fit the ellipse to the Chebyshev polynomial. The Chebyshev polynomial is regarded as the amplification factor of a multistage scheme.

The application of these optimized schemes to actual Navier-Stokes calculations awaits execution. This work is done in collaboration with Bill Kleb (NASA Langley).

KUN XU

Gas-kinetic Theory Based Flux Splitting Method for Ideal Magnetohydrodynamics

The construction of approximate Riemann solvers for the ideal MHD equations is very important to the understanding of global heliosphere. However, even for the 1D MHD equations, there are complicated eigen-systems. For example, entropy, slow, Alfvén, fast waves all have to be considered in the development of a single flux function in the approximate Riemann solver. Due to the nonstrict hyperbolicity of the

MHD equations, there is the possibility that two or more eigenvalues of the Jacobian matrix are equal. In some cases there could be infinite eigen-directions in phase space, which leads to the possible breakdown of the waves structure. The present work is related to the development of a gas-kinetic scheme for the MHD equations, where the complicated eigen-structure of MHD waves are totally abandoned.

The gas-kinetic MHD solver is based on the splitting of the macroscopic MHD fluxes directly. In the gas-kinetic scheme, all flow transport is associated with the particles moving across cell interfaces. For example, we can easily decompose the mass density into two groups. One is moving to the right and another is moving to the left. Similarly, the momentum and energy densities can also be split according to gas-kinetic theory. In generalization, we can use the same techniques to split the corresponding quantities, such as pressure, energy fluxes, etc., in the MHD equations. For magnetic field, it is considered to be frozen into the particle motion. As a passive scalar, it is moving along with the particle transport. The final fluxes at a cell interface are composed of the fluxes from the left side with positive particle velocity and that from right side with negative velocity. At the same time, the free transport mechanism is modified by the particle collisions at the cell interface in order to reduce the artificial dissipation. We have tested the above scheme and compared its numerical results with the well-developed Roe-type MHD solver. For the Brio-Wu test case, the simulation results are almost identical from both approaches, but the new kinetic method is three times cheaper than the Roe MHD solver. In high speed MHD flow simulations, the kinetic scheme is even more robust. The work performed thus far indicates that the approach taken is sound and there is enormous economic value to continue the current research.

We plan to apply the new scheme to the study of global heliosphere and solar wind heating problem for the understanding of the living environment of the earth.

AMIR YEFET

Construction of Three Dimensional Solutions for the Maxwell Equations in Close Domain

We consider a family of solutions for the three dimensional Maxwell equations in free space and in cubic cavity.

We used these solutions as test cases for numerical solutions for the three dimensional Maxwell equations in free space.

We have completed coding the Yee and Ty(2,4) schemes in FORTRAN 77 and have begun the debugging phase.

Future plans are to use the Ty(2,4) scheme for practical uses.

This work was done in collaboration with Karl J. Moeller (NASA Langley) and Prof. E. Turkel.

A Note on Numerical Boundary Conditions for the Ty(2,4) Scheme

Numerical simulations of the finite difference Ty(2,4) scheme show that when a mesh is refined the order of the Ty(2,4) scheme decreases. This can be the result of an accumulation of round-off errors or the use of insufficiently accurate computers. Here we suggest the use of different numerical boundary conditions, which allow us to overcome this limitation.

Future plans are to use these new numerical boundary conditions for other examples and to prove that these numerical boundary conditions are stable.

The Effect of Absorbing Boundary Conditions on the Behavior of Non-periodic Solutions

Most of the absorbing boundary conditions, which are based on Berenger PML technique, assume periodic solutions. We examined four techniques using the Ty(2,4) scheme both on periodic and non-periodic waves which, after a short period of time live the computational domain. We look at the electric and the magnetic field at this time after using absorbing sponge layers.

The Maxwell equations are solved by the Ty(2,4) scheme inside the computational domain. The computational domain is surrounded by an absorbing layer which is an aggregate of PML media, which satisfy either Berenger PML technique equations or the Lorentz material model equations by Zilkowsky or Abarbanel and Gottlieb PML equations or the UPML equations by Gedney. At the outer edge of the PML domain we impose perfectly conducting conditions.

Tests show that in three of these schemes the error in L_2 norm increases after a long period of time.

Future plans are to use modifications done by S. Abarbanel and D. Gottlieb in order to overcome this anomaly. This work was done in collaboration with Prof. S. Abarbanel.

PHYSICAL SCIENCES, FLUID MECHANICS

P. BALAKUMAR

Periodic Secondary Motions in Core-annular Flows

Our objective is to compute the nonlinear equilibrium motions in core-annular flows. In this method, the disturbances are assumed to be periodic in the flow direction and to be steady in a frame moving with a constant phase velocity. In the next step, the disturbances are expanded in Fourier series in the flow direction and the wavenumber, the phase speed and the Fourier components are solved for directly in an iterative manner by solving the full Navier-Stokes equations. The main difficulty in applying this method to problems with interfaces is that the shape of the interfaces is not known a priori and have to be determined as part of the solution by satisfying the jump conditions and the kinematic conditions across the interfaces. To overcome this problem, we first make the interfaces and the walls parallel to each other by transforming the normal coordinate and we use the Fast Fourier Transform (FFT) method to find the forcing terms for each Fourier components.

Starting with the linear neutral solutions, nonlinear neutral curves are mapped out by gradually increasing the magnitude of the disturbances. The nonlinear neutral curves associated with the viscosity differences induced instability from a closed curve starting from one linear neutral solution and ending in the other neutral solution. The nonlinear neutral curves associated with capillary instability with large viscosity ratios formed open curves. The interface dips towards the center and oil bubbles form. The neutral curves with small viscosity ratios formed closed curves and the waves appear as bamboo waves.

In the future, we plan to compute the nonlinear neutral solutions and the free surface shapes in flows with free surfaces, e.g., free falling film and water waves.

CHING-CHANG CHIENG

Flow Control

The objective of this research was to set up a computational tool and simulate the flow phenomena of oscillatory blowing.

The NTHU (National Tsing-Hua University) computer code was revised during our three-week visit to be able to use the ICASE computers to perform the computation simulating the flow field of oscillating blowing. It was found that the flow control problems are very intensive in computation. The ICASE computers are somewhat inadequate in speed for the flow control studies. A basic case was completed using the NACA-15 model simulating the flow field at Mach 0.55 with an oscillatory jet at the frequency 775Hz. The results indicated that the NTHU computer code worked well and could be used in future studies for comparison with experimental data.

The flow control studies continue to progress at NTHU. A paper is being prepared for submission to AIAA Applied Aerodynamics Conference to be held in the Norfolk Convention Center, June 1999.

This work was done in collaboration with Jer-Nan Juang (NASA Langley).

SANG-HYON CHU

Development of Microwave-driven Smart Material Actuator

“Wireless” actuators controlled remotely with microwaves offer tremendous advantages over hard-wired actuators, especially for space applications such as the Next Generation Space Telescope (NGST), in which thousands of discrete actuators are required to effect high precision distributed shape-control of the primary reflector. This new concept alleviates the need for hard-wire connections resulting in significantly simpler system designs and lower system mass. “Wireless” control of an actuator involves constructing actuator elements with layers of metal conductors embedded into smart material. The H-field of microwave that is incident on a conductor instantaneously induces a current or an electric field between metal conductors. The desired motion or displacement of actuators is, therefore, made possible by controlling an electric potential due to the coupled microwave energy.

3x3 rectenna patches built at JPL were received and tested in an anechoic chamber by modulating microwave power level, frequency, incident angle, and polarization angle. It could produce up to 73 Vdc. PZT 5A multilayer piezoelectric actuator was selected as the smart actuator and tested with an independent power source before a direct coupling with 3x3 rectenna was built up. The experimental results of the direct coupling showed that rectenna patch could produce enough voltage to activate smart actuator by the correction limit suggested by NGST program. According to test results, the voltage variation by various factors allows much more and wider spectrum of controllability for displacement of actuators.

In the future, it is planned that an array of rectenna/actuator patches will be developed and tested for the threshold requirements of smart material actuators. The networked array concept of control logic-embedded microwave-driven smart actuators will also be studied and applied to this project.

WILLIAM O. CRIMINALE

Receptivity and Transition in Boundary Layers

Flow control has been and remains an essential goal in fluid mechanics. In many cases, control means retarding any transition or break down. At other times this may mean enhancing the action so that mixing might be made more efficient, for example. In this respect the laminar boundary is one of the most salient prototypical examples in all of fluid mechanics. Unlike many other flows, the boundary must interact with a free stream that is known to affect the dynamics. Then, the major cause for instability of the flow is due to the viscosity that created it. This combination makes for a formidable problem. Nevertheless, the goal is to control the flow in such a way that drag reduction becomes possible. Such a task requires an understanding of the transient dynamics. Simple computations to ascertain the stability is not the correct premise for this kind of exploration.

This investigation uses three means for examining the problem. First, analytical techniques are used to investigate the complete linear perturbation equations with the possibility of both vorticity and pressure in the free stream. This is done as an initial-value problem. Second, full numerical treatment of the linear system is made to verify the analytical predictions. Finally, a direct numerical simulation for the boundary layer is made using the bases derived from the smaller amplitude predictions. It is believed that this work is the first based on principles that actually reveals how receptivity actually works.

One work describing the early period dynamics for the perturbations in the boundary layer will appear in the *Journal of Fluid Mechanics*. A second report will be available soon and deals with the subsequent

dynamics and transition. Future work will concentrate on two more aspects of this important problem, namely (1) spatial evolution of disturbances and (2) more innovative analytical means for solving the linear perturbation system.

This work was done in collaboration with D.G. Lasseigne, T.L. Jackson, and R.D. Joslin.

AYODEJI DEMUREN

Computations of Turbulent Buoyant Flows

In the numerical computations of turbulent buoyant flows there are buoyancy effects on the mean flow and turbulent parameters. It is usual to model these effects with the evocation of the Boussinesq approximation of the buoyancy force in the mean flow equations, which is then carried on to the gravitational production terms in the turbulent flow equations. This approximation has limited applicability and should be avoided in totality in general computer codes. Furthermore, the gravitational production terms are directionally sensitive and current models artificially change coefficient values depending on the flow orientation with respect to gravity. This practice is unsatisfactory for general flow computations. The goal of the research study was to eliminate these restrictions, and derive a generalized treatment suitable for the computation of all types of turbulent buoyant flows, irrespective of density difference or flow orientation.

Equations are derived which show the limit of applicability of the Boussinesq approximation; for ideal gases this translates to the requirement that temperature differences should be much lower than the reference temperature. Alternative formulations are presented to overcome this limitation and be equally accurate and robust for all possible ranges of temperature and density differences. Effects of buoyancy production on the turbulent kinetic energy are direct and not controversial, but the effect on its rate of dissipation has to be made directionally sensitive to reproduce experimental observations. A generalized treatment is derived which has the desired effect in all flow inclinations, with respect to gravity, without the need to modify any coefficients. Comparison of model predictions to experimental data for strongly heated ducts in natural and forced convection show good agreement.

Advanced models for the effects of buoyancy on the production/destruction of dissipation are to be considered. This work was done in collaboration with Bob Rubinstein.

REMI DRAI

Pole Placement in Sectors via Static Output Feedback

The problem of the design of a static output feedback gain for stabilizing a linear time invariant dynamical system can be fruitfully addressed in the framework of convex optimization by solving a sequence of semi-definite programs. For most applications however, stability is not the only requirement so the usual Lyapunov techniques need to be extended for handling other design specifications. It is well known, for example, that the rapidity of the transient behavior of a system can be obtained by considering shifted complex half-planes instead of the classical stability left half-plane. The concept of LMI regions and the corresponding stability criteria given by Gahinet-Chilali can be used for further extensions. For example, placing the eigenvalues of the closed loop system in an appropriate sector region of the complex left half-plane may insure satisfactory damping and overshoot properties.

In this work, we give sufficient LMI conditions, that when satisfied, achieve the above program. An extension of the classical stability result of Lyapunov is used to express the pole constraints, resulting in a

nonlinear hermitian matrix inequality. After performing a completion of square procedure, one nonlinear term is eliminated by introducing a supplementary matrix variable. An application of the Schur's lemma as well as some definiteness properties of the global constraint result in a set of two real linear matrix inequalities. An algorithm intended to reduce the sufficient and non-necessary character of the conditions is then provided. Moreover, additional constraints on the static output feedback gain, such as a block-diagonal structure or positive definiteness, can be readily achieved.

Extension of this work to the case of uncertain (polytopic) models will be considered in the future and numerical examples provided. Helpful discussions with S.M. Joshi (NASA Langley) are gratefully acknowledged.

QIANG FU

Improvements of Fu-Liou Radiation Model for the CERES Applications

Clouds and the Earth's Radiant System (CERES) is a research project led by the Radiation Sciences Branch, NASA Langley Research Center, to examine the role of cloud/radiation feedback in the Earth's climate system. In this project, the Fu-Liou radiation model is used to derive the shortwave and longwave radiation fluxes at the surface and within the atmosphere. The objective of the present research is to improve the model for broadband longwave fluxes and develop a sound working code for the unfiltered radiance in the CERES window channel.

Improvements to the broadband longwave fluxes have been made by

1. including minor gaseous absorption bands due to CO_2 and CFCs in the window and O_3 outside $980\text{-}1100\text{ cm}^{-1}$
2. slight increase in number of k 's used
3. extending the H_2O continuum to cover the entire longwave range

We have developed a working code for both broadband and window radiances, which considers the multiple-scattering effects. The two-step procedures have been successfully formulated to convert the Fu-Liou window unfiltered radiance ($800\text{-}1250\text{ cm}^{-1}$) to the CERES unfiltered window radiance ($847\text{-}1219\text{ cm}^{-1}$). A newly developed cirrus radiative properties parameterization has been incorporated to the radiation model, which affects both shortwave and longwave fluxes and requires a consistent treatment of cirrus in the CERES cloud retrieval and the modified Fu-Liou model in SARB. The work performed thus far indicates that the modified Fu-Liou model is well-suited for the CERES longwave SARB application.

Some improvements need to be done in the shortwave. Future plans have been worked out.

SHARATH S. GIRIMAJI

Pressure-strain Correlation Modeling: Testing and Validation

Pressure-strain correlation is one of the key processes that need to be modeled for accurate prediction of turbulence. The objective of this study is to develop and validate such models for complex turbulent flows.

The pressure-strain correlation model developed earlier is now entering the testing and validation stage. It is being tested in two flows dominated by rotation. (i) The first is a laboratory flow called the rotating channel flow. In this case, we are interested in determining whether the model is capable of reproducing the stabilizing and destabilizing effects of rotation. (ii) The second test flow is the more realistic trailing vortex in an aircraft wake. Here, the model is being used to examine the various decay mechanisms of trailing vortices.

The testing is in its preliminary stages and validation over a wider range of flows will come next.

This research was conducted in collaboration with S. Thangam (Stevens Institute) and S. Wallin (FFA, Sweden).

Rotating Turbulent Flows

In many aeronautical flows, turbulence is subjected to rotation. The objective of this fundamental study of great practical importance is to better understand the behavior of turbulent fluctuations in rotating flows.

The first stage of this study is complete, wherein a generalization to the Taylor-Proudman theorem was proposed. As a second step, I am now investigating the exact solutions of linearized perturbation equations of rapidly rotating two-dimensional flows. These solutions can shed more light on the validity of the previously proposed theorem and further enhance our fundamental understanding. In particular, there appears to be an important difference in the manner in which vortical and irrotational fluctuations behave. This issue is currently under investigation.

Close examination of other basic flows subject to rotation is the next step.

This work was done in collaboration with J.R. Ristorcelli (formerly of ICASE).

Non-equilibrium Algebraic Reynolds Stress Modeling

The full Reynolds stress closure is computationally too prohibitive for many practical applications. In this study, we plan to devise simplified versions of this model (algebraic Reynolds stress model) which are computationally more viable.

A general theory for reducing the degrees of freedom of non-linear systems is under development. This is based on the reduction of overall evolution potential of the equation system. This theory will immediately lead to better non-equilibrium algebraic Reynolds stress models. The theory will also provide a theoretical foundation for the previous algebraic models.

Completion of theoretical development followed by validation and testing will come next.

One-equation Turbulence Models for Acoustic Calculations

Accurate calculations of the flow field surrounding an aircraft is necessary for computing the sound generated by the aircraft. One of the key turbulence quantities required for sound estimation is the length scale of the flow field. The currently used one-equation model provides little information about this quantity. Our objective is to develop more advanced one-equation models which provide the required information.

I am in the process of deriving a one-equation turbulence model by subjecting the Reynolds stress closure equations to plausible simplifying assumptions. It is expected this model will be more accurate than the current model and provide the all-important length scale information as well.

In the near future, I expect to complete the model development.

C.E. GROSCH

Simulation of Supersonic Jet Mixing in Lobe Ejectors

Mixing enhancement of high and low speed streams is utilized as a means to improve efficiency of supersonic combustors, reduce aircraft signatures, and control high speed jet noise. One common method of mixing enhancement is to use lobe mixer ejectors. Another is to place tabs on the edges of the jets. In the main, experimental studies are available to evaluate the performance and guide the design of these mixers.

The objective of this research is to use numerical simulation to examine the performance of lobe ejectors, with and without tabs, in order to understand the physics of the mixing and how it is affected by changes in the parameters of these devices.

A set of numerical calculations were carried out using the compressible three-dimensional Navier-Stokes equations. The tabs were modeled by pairs of counter rotating vortices. Various geometric configurations of the lobe mixers are being used together with periodic side boundary conditions to simulate an array of these devices.

Simulations with tabs used an axisymmetric round inflow for the jet with no tabs, two tabs and four tabs arranged symmetrically on the edge of the jet. The results of the simulations show a striking similarity in the flow evolution between numerical and experimental jets, particularly in the jet potential core region. The experimental and computational results also show similar values for mass entrainment. Because of complications associated with computational boundaries, the comparison between experiment and simulation degrades beyond 8 jet diameters downstream. These results were presented by my co-author, J.M. Seiner, at the Fourth AIAA/CEAS Aeroacoustics Conference in Toulouse, France, 2-4 June 1998.

The simulations of the lobe mixer flow show that the jet becomes unstable and oscillates in the "garden hose" mode. For a particular lobe geometry and velocity ratio, the oscillation has a constant, narrow band, frequency near the inflow. Further downstream the amplitude grows and the motion becomes non-linear leading to spectral broadening. Typical Strouhal numbers of the narrow band oscillation is about 0.45. The physics of this phenomena is related to the no slip boundary condition at the top and bottom walls which cause large shear, generating streamwise vorticity. As the disturbances become nonlinear the streamwise vorticity migrates from the solid boundaries into the interior. Strong mixing follows and, by about half way down the channel, the jet and coflow become nearly fully mixed.

Future experimental and numerical studies are required to more clearly define the initial induced vorticity field in the round jet. Future experiments will use PIV imaging to measure the cross-stream vectors. The experimental data will be used to set the inflow conditions for the numerical simulations.

Further calculations are planned for the lobe mixer including varying the geometry and adding tabs on the sides of the lobes.

G.M. LILLEY

The Study of Airframe Noise and its Reduction and the Study of Jet Noise for the Evaluation of its Spectrum at High Frequencies and its Directivity as a Result of Refraction

The control of flyover noise from an aircraft in its high-lift, gear down configuration has received renewed attention as a result of widespread complaints from people living near airports and the need for all aircraft to meet the legal limits of aircraft noise as specified by the present ICAO and FAA Stage 3 limits and the expected new lower Stage 4 limits. Aircraft noise prediction demands a detailed evaluation of both the steady and unsteady flow over the entire aeroplane. The demand is for CFD calculations which are time accurate and have high spatial resolution. Small changes in the flow structure may be relatively unimportant in steady flow predictions but can be intense sources of noise. Studies of airframe noise demand both high quality experiments in large facilities as well as numerical calculations for the unsteady flow field and their matching with calculations for the radiated farfield noise. The numerical methods used for the unsteady flow calculations need to be very carefully calibrated against exact analytical solutions since large errors are possible, for the radiated noise is only a vanishingly small percentage of the flow field fluctuations.

In recent years the study of numerically based methods has only just begun to have any reliability and provide predictions of noise of given aircraft to the accuracy demanded by Noise Certification. In the past all prediction methods have necessarily been based on accumulated experimental model and full-scale data bases with very little input from theoretical methods apart from dimensional analysis and source identification methods. Noise reduction has largely been the result of ad hoc experiments.

The objective has therefore been to study (a) time accurate numerical and analytical methods to provide the unsteady flow field database needed in the Lighthill Acoustic Analogy to predict the farfield radiated noise, its spectrum and directivity, and (b) to study methods of noise reduction both for the airframe and the engine. The reduction of airframe noise clearly has a parallel with the noise generation of all classes of birds and the silent flight of the 'owl.' Now the linear scale and flight speed of the 'owl' are minute compared with an aircraft and the flight Reynolds number is well below one million. Nevertheless, the 'owl' generates almost zero noise at frequencies greater than 2kHz at high flight lift coefficients comparable with that of an aeroplane. This has only been achieved by delicate modifications by Nature to its feathers causing complex changes to the flow field over its wings. A detailed understanding of these (favorable) flow changes appeared worthy of further study to ensure that the same 'technology' was not overlooked when searching for methods of airframe noise reduction.

The development of a flow solver for the steady and unsteady flow over components of an aircraft and jet has been in progress for about the past three years. At Southampton University in UK, Zhang, Rona and Lilley have concentrated on improvements to a time-dependent RANS solver, called 'TRANS,' which is fast and was especially developed for supersonic self-excited flows, such as the rectangular cavity and the shock cell noise of a jet. Part of my time has been spent checking our two-equation turbulence model against that used in CFL3D and to prepare our paper for publication. Many doubts have been raised by researchers in the U.S. regarding this methodology and its application to flows at low subsonic Mach number and, especially to flows that are not in the class of self-excited flows.

There is now evidence that the 'TRANS' solver, with carefully chosen grids in flow regions of great complexity, and covering a wide range of Reynolds number, is capable of providing accurate unsteady flow input data into the Ffowcs Williams and Hawkings form of Lighthill's Acoustic Analogy from which the farfield radiated noise can be calculated. At low subsonic Mach numbers, the effects of refraction can be neglected but the method allows for the radiated sound to be calculated even when suffering reflection and scattering from complex surfaces. The calibration of the 'TRANS' methodology for low Mach number subsonic flows has been made by comparison with good experimental data and by comparison with simple unsteady flows for which analytic solutions were available.

The further study of 'owl' feathers was made possible by a visit to Dr. K. McKeever, Director of the "Owl Association," at Vineland, Ontario, Canada. Dr. McKeever was able to explain the different hunting habits of the various owls in Canada and the U.S., and which species were dependent on 'silent flight' for catching their prey. I was able to obtain a large number of feathers from Dr. McKeever which I am in the process of studying. I have obtained good magnification of the leading edge comb on the primary feathers and the trailing edge fringe. Previous work had shown the 'leading edge comb' was responsible for providing attached flow over the owl's wing at a lift coefficient of about unity, and thus allowed the owl to fly unstalled in both flapping and gliding stable flight towards its prey. The noise was also reduced since the attached flow promoted a thinner boundary layer at its trailing edge than the normal boundary layer in an adverse pressure gradient. The purpose of the 'fringe' at the trailing edge was to provide a pressure release and thus prevented trailing edge noise scattering, a source of increased radiated noise.

The study of flow-noise interaction has been continued and a paper has been prepared reviewing the different approaches made during the past 25 years. These methods have been compared with the Lighthill Acoustic Analogy. The paper will appear in a monograph in preparation on aerodynamic noise. A further part of this monograph has been devoted to an explanation of the comparison between methods used in acoustics for sources traveling at speeds comparable with the speed of sound and electromagnetic sources of electromagnetic waves traveling at speeds near that of light. In this work, I have made contact with Dr. Thornhill in UK who has also worked on this problem.

The flow calculations in 'TRANS' are, in general, under-resolved at high frequencies arising from the use of a turbulence model and the limitations made regarding the grid dimensions for a high Reynolds number flow. It would be useful therefore to investigate the frequency decay law that would be expected at high Reynolds numbers. Work by Tam and Seiner (1997) on collating a vast quantity of experimental data on the spectrum of the radiated noise from jets at subsonic and supersonic speeds has shown that, at least for jet noise, the high frequency portion of the noise spectrum has an almost universal character. An attempt has been made to find the high frequency decay law for isotropic turbulence at low Mach numbers. Guided by the results obtained by Dubois (1993) and Sarkar and Hussaini (1993) for the retarded time-space covariance at low Reynolds numbers in isotropic turbulence, it is found that in the range of acoustic frequencies beyond the peak in the radiated noise spectrum, including the inertial sub-range, a universal spectrum, $\omega^{-3/2}$ is obtained similar in form to that derived by Zakharov (1965). It is shown that this result is flow-dependent and that for the radiated noise from a self-preserving turbulent shear layer the law would be ω^{-2} in this same restricted frequency range. The jet noise data of Tam and Seiner (1997) show from measurement the decay law ω^{-n} , where n varies between 2.2 and 2.5.

LI-SHI LUO

Lattice Boltzmann Scheme for Flow-structure Interaction

One important problem in the application of the lattice Boltzmann equation to various flow problems is the interaction between fluid flow and solid boundaries, i.e., the implementation of boundary conditions in fluid-structure interfaces. Moving boundary problems in high Reynolds flow pose a challenge to traditional CFD methods. Usually, turbulence modeling has to be employed in such cases. The present work uses the method of the lattice Boltzmann equation (LBE) to simulate a flow-structure interaction problem.

With the LBE method, boundary conditions for objects with complicated geometries are easily implemented. We intend to implement a computationally efficient boundary condition for moving boundaries in high Reynolds number flows. Various schemes combining existing bounce-back type boundary conditions with interpolation (or extrapolation) are under theoretical study and numerical test.

The present work has been funded by NASA Langley Research Center under the program of "Innovative Algorithms for Aerospace Engineering Analysis and Optimization." The Co-PI's of the proposal for the present work are Prof. Renwei Mei (UFL), Dr. Li-Shi Luo, and Prof. Wei Shyy (UFL). The collaboration also includes Prof. Pierre Lallemand (Director, ASCI-CNRS, Univ. Paris-Sud), and Prof. Dominique d'Humières (ENS, Paris).

Lattice Boltzmann Model for Non-ideal Gases

The key issues in the study of multi-phase (e.g., liquid-vapor) flows are the modeling of interfaces and phase transition among different phases. It is difficult to use the Navier-Stokes equations to model the

inhomogeneous multi-phase flows because interface tracking is a laborious computation. In the past few years, a number of lattice Boltzmann models have been developed to model multi-phase flows. However, the multi-phase lattice Boltzmann equation is still lacking a rigorous theoretical basis. For instance, previous multi-phase lattice Boltzmann models do not have a consistent equilibrium thermodynamics. The present work applies the Enskog theory of hard spheres to revise to theory of the multi-phase lattice Boltzmann equation.

With the Enskog theory, we were able to derive a new multi-phase lattice Boltzmann model which has consistent equilibrium thermodynamics. We have rigorously demonstrated the deficiencies in the previous multi-phase lattice Boltzmann models and provided a systematic procedure to derive a correct multi-phase lattice Boltzmann model based upon the Enskog theory (or the revised Enskog theory). A brief account of the present work has been published in Physical Review Letters and as an ICASE report. An extended version of the work has been submitted to Physical Review E and a corresponding ICASE report is in preparation.

We intend to derive a thermodynamically consistent multi-component lattice Boltzmann model in the future based upon the same methodology.

ALEXANDER MURAVYOV

Application of Equivalent Statistical Linearization Methods to the Problem of Nonlinear Random Vibrations

Design of high speed vehicles necessitates the further development of sonic fatigue technology to understand the fatigue mechanisms and to estimate the service life of aerospace structures subjected to intense acoustic and thermal loads. Further improvements in vehicle performance and system design are hampered by the limited understanding of the physical nature of geometrically nonlinear structural response. Conventional (linear) prediction techniques lead to grossly conservative designs and provide little understanding of the nonlinear behavior. The objective of the current research includes development of equivalent linearization methods for the problem of random vibrations of geometrically nonlinear structures excited by white noise. Dynamics of such structures is described by stochastic differential equations.

The research is currently done on application and development of equivalent linearization methods for solution of stochastic differential equations. Particularly, a new version of the equivalent linearization method based on the minimization of the strain energy error is developed. The preliminary results show that this method yields more accurate results than the methods based on conventional (force error) minimization.

The future plans include the further development of the equivalent linearization technique based on the minimization of mean-squared error in strain energy and application of this method to a wide class of finite element nonlinear models with developing an in-house code which is designed to work along with the MSC/NASTRAN code for application to random nonlinear vibrations.

This work is done in collaboration with S. Rizzi, T. Turner, and J. Robinson (NASA Langley).

ALEX POVITSKY

Computation of Three-dimensional Acoustic Fields

Direct solution of narrow band linear systems (Thomas algorithm) on parallel computers is essential for compact high-order methods, implicit numerical schemes and alternating line methods in multi-grid.

Known ways to parallelize the Thomas algorithm include use of numerical boundary conditions between subdomains belonging to different processors and use of factorization algorithms. The former methods eliminate far-field data dependencies; however, modification of either the finite-difference approximation or implicitness of the scheme due to interface boundary conditions can cause deterioration of the accuracy, stability and convergence properties relative to the original serial method. The latter methods increase the computational work by more than a factor of two. We propose a way to reduce the parallelization penalty of the Thomas algorithm where the numerical algorithm remains the same as the serial one and only the order of computations is changed.

Our approach is based on the consideration of solution of multiply banded systems as part of an entire solver which typically includes solutions of such systems in three spatial directions, local computations of coefficients of discretized non-linear Euler/Navier-Stokes equations and Runge-Kutta temporal computations. Following are reasons for poor parallelization efficiency of the parallel pipelined Thomas algorithm: (i) there is no complete data for other computational tasks while processors stay idle and (ii) communications control computational tasks as either the forward step coefficients or the backward step solution must be obtained from the neighboring processors for the beginning of the forward or the backward step computations. To overcome the first problem, a new pipelined Thomas algorithm named the Immediate Backward Pipelined Thomas Algorithm (IB-PTA) has been developed at the previous stage of this research. To overcome the second problem, a new scheduling algorithm is proposed. The advantage of the “control by schedule” over the “control by communications” is that processors do not wait for data from other processors. Instead, processors compute other tasks and switch to receive data only when these data are available on neighboring processors and necessary in a current processor. For the parallel version of the ADI code FDL3DI (developed by the Wright-Patterson Air Force Base) the idle processors are used for data-independent computations of discretized coefficients and local data-dependent multiplications of intermediate ADI functions by transformation metrics. The measured parallelization penalty on 64 processors of Cray T3E is varied from 18% on 10^6 nodes to 22% on 2.16×10^5 nodes comparative to 33% – 99% for the basic pipelined algorithm. For compact solvers banded linear systems are solved independently in three spatial directions; therefore, the idle processor time is used for solution of these systems in the next spatial direction. Runge-Kutta computations are scheduled while processors are idle in the last spatial direction. Here we use the IB-PTA to ensure that spatial derivatives in this direction have been computed by this time.

We plan to parallelize compact aeroacoustic codes of LaRC (Mark Carpenter) and WP-AFB (Miguel Visbal) by the proposed method. We will implement this approach to parallel load-balanced multi-block computations where each processor will be able to perform computations originated from the other numerical grid while it stays idle from the Thomas algorithm computations on the current grid.

ROBERT RUBINSTEIN

Shock Wave Propagation in Weakly Ionized Gases

Recent proposals for a high-speed civil transport have revived interest in the possibility, suggested by some experiments performed about a decade ago in the former Soviet Union, that shock waves can be mitigated by the addition of charged particles to the flow. The goal of this investigation is to contribute to the theoretical understanding of what mechanisms might be responsible for these observations.

This work has focused on the role of the energy source ahead of the shock. One possibility is that it moves the energy distribution away from the Maxwellian, maintaining a steady state far from equilibrium.

An analogous problem is the shock in a vibrationally excited medium, in which it is known that the dynamics can be significantly altered by the presence of a nonequilibrium state on one side of the shock. Modified shock jump conditions for a model problem which reflect the existence of a far from equilibrium distribution on one side of the shock have been analyzed. It is found that this arrangement increases the shock speed. The question remains open whether the energy requirements are less than those of direct heating.

This methodology will be generalized to a system containing charged particles to better simulate the properties of a weakly ionized gas.

This work was done in collaboration with A.H. Auslender (NASA Langley).

Helicity Effects on Turbulent Sound Radiation

The addition of swirl to turbulent jets had been proposed as a means to reduce jet noise. Swirl was found not to reduce total sound levels, but it will be useful to understand what effects should be anticipated from swirl addition.

Following our earlier work on sound radiation by isotropic turbulence, the effect of helicity on turbulent time correlations has been analyzed. It is found that the Eulerian time correlations relevant to sound radiation are strongly modified by helicity, and that the correlation times are increased. The effect should be a redistribution of acoustic energy to lower frequencies.

Comparisons with experimental data from the Langley Jet Noise Laboratory are in progress.

This work was performed in collaboration with Ye Zhou (ICASE and IBM).

Theory of Rotating and Stratified Turbulence

Heuristic scaling laws proposed for rotating turbulence have proven useful in the derivation of turbulence models for rotating flows. Some unresolved issues concerning these proposed scaling laws are addressed in this work.

An open problem in the weak turbulence theory of rotating turbulence, the effects of possible double resonances, has been addressed. The multiple resonances have been identified and their effect on the scaling laws has been derived. It is found that there are small corrections to the scaling laws, but that these are unlikely to invalidate the turbulence models.

A second problem, the demonstration of the locality of the spectrum, or its independence of the large scale excitation and small scale dissipation mechanisms has also been analyzed. It is found that rotating turbulence is only marginally local, and that the large scale excitation does exert a weak influence on the scaling law. Again, no significant modification of existing models is required.

This work will be extended to the analysis of the dissipation range in rotating turbulence.

CHI-WANG SHU

High-order Discontinuous Galerkin Method and WENO Schemes

Our motivation is to have truly high-order methods for structured and unstructured mesh which are easy to implement for parallel machines. The objective is to study and apply high-order discontinuous Galerkin finite element methods and weighted ENO schemes for convection dominated problems. The applications will be problems in aeroacoustics and other time dependent problems with complicated solution structure.

Jointly with Harold Atkins at NASA Langley, we are continuing in the investigation of using the discontinuous Galerkin method to solve the convection dominated convection diffusion equations. Emphasis is

put upon studying the parallel implementation issues of discontinuous Galerkin methods for viscous terms. Efficient parallel implementation requires certain modification to the viscous fluxes, the effect of which upon stability is being studied. Jointly with Changqing Hu (Brown University), we have designed a class of third- and fourth-order weighted ENO schemes for unstructured mesh. Adaptive methods based on such WENO schemes are being studied.

Research will be continued for high-order discontinuous Galerkin methods and weighted ENO methods and their applications.

SIVA THANGAM

Development and Analysis of Turbulence Models for Complex Flows with Curvature and Rotation

Significant levels of curvature and/or rotation is often present in flows of practical importance. Turbulence models for such flows are required to accurately depict the flow field and turbulence stresses. The focus of research during the current phase involves the development of efficient turbulence models ranging from anisotropic two-equation models to full second-order Reynolds stress closure schemes that can capture the effects of curvature and rotation as well as their validation for various benchmark test cases. The work performed during the current period involves my collaborative efforts with Sharath Girimaji and Ye Zhou (ICASE) as well as others at Stevens Institute of Technology.

Two-equation turbulence models for rotating flows have been developed based on a) the phenomenological treatment of rotation modified energy spectrum and b) generalized algebraic representation of the second-order closure. Such models account for the influence of solid body rotation in the energy transfer process. The model development includes the development of a suitable dissipation rate equation with rotation rate dependent model coefficients. The model has been successfully applied to flows in rotating duct and wake flows, among others.

Investigations on the development and implementation of a generalized Reynolds-stress closure scheme for complex flows has recently been started in collaboration with Girimaji. The scheme was first developed by Girimaji (1998) to include proper limiting behavior as well as both the imposed and frame rotation and was validated for various homogeneous flows. The model has recently been extended to rotating channel flows and validated against available LES and experimental data.

Future work will involve the application of the two-equation and the second-order closure models to complex flows.

STEFAN WALLIN

Numerical Study of Far Field Turbulent Trailing Vortices

The long time decay of aircraft trailing vortices at take-off and landing has been considered. The importance of the vortex self-generated turbulence vs. ambient influences as the vortex decay mechanism has been studied by doing RANS computations using state-of-the-art Reynolds stress transport (RST) modeling. Moreover, the importance of different modeling approaches has been assessed by comparing different published RST models with RST models developed at ICASE and algebraic Reynolds stress models developed at ICASE and FFA.

A code was written specifically for developing vortices with the aid of the MapleV software. The governing equations and boundary conditions were given on differential form and a MapleV script written for this purpose automatically discretizes the equations and generates Fortran codes that solve the problem.

The work was done in collaboration with Sharath Girimaji (ICASE) and Fred Proctor (NASA Langley) and will result in two reports (and articles).

YE ZHOU

Turbulent Force as a Diffusive Field with Vortical Sources

In Reynolds-average Navier-Stokes equation, the divergence of Reynolds stress tensor, i.e., the turbulent force, rather than the tensor itself, is to be simulated and partially modeled. Thus, directly working on turbulent force could bring significant simplification.

A novel exact equation for incompressible turbulent force is derived. All source terms to be modeled are vortical, the dominant mechanism being the advection and stretching (with an opposite sign) of a “pseudo-Lamb vector” by fluctuating velocity field. No coupling with pressure is involved. The equation follows from a study of the mean fluctuating Lamb vector and kinetic energy, which constitute the turbulent force. Both constituents are governed by the same kind of equations. This innovative equation is similar to Lighthill’s acoustic analogy and naturally calls one’s attention to studying the vortical sources of turbulent force. The methodology described here may lead to turbulence models which provide more complete treatment than that of two-equation models, but relatively easier computation than that of second-order closures.

We plan to implement this new modeling scheme in realistic scientific and engineering flows.

This research was conducted with Drs. J.Z. Wu and M. Fan (University of Tennessee Space Institute).

COMPUTER SCIENCE

SHAHID H. BOKHARI

An Evaluation of the Tera Multithreaded Architecture for Unstructured Meshes

The Tera Multithreaded Architecture (MTA) is a new parallel supercomputer currently being installed at San Diego Supercomputing Center (SDSC). The machine uses custom hardware to implement a number of innovative ideas, such as hardware support for multithreading, very fine-grained synchronization and a flat shared memory without locality. The objective of this research was to evaluate the performance of the Tera on unstructured mesh problems, which are difficult to parallelize efficiently on conventional parallel computers.

A representative kernel from Dimitri Mavriplis's Eul3D code for solving the Euler equation on an unstructured mesh was ported to the Tera. The original Fortran code required very little modification, and parallel execution was measured by varying the number of streams (i.e., active threads) and the number of processors (currently only one or two). The machine exhibits excellent speedup as the number of streams is varied, however when the number of processors is changed from one to two, the performance improvement is limited by the present incomplete interconnection network. Overall, we were able to achieve more than 200 MFlop/s per processor, which is quite respectable.

We are looking forward to repeating our measurements on larger Tera configurations. The Tera at SDSC is slated for upgrading to four processors very soon and to eight processors by next summer. We plan to rerun our experiments to obtain a clear estimate of the scalability of the architecture. We also wish to look at a more comprehensive set of experimental measurements involving meshes of varying sizes. A multigrid version of Eul3D will also be implemented.

This work was done in collaboration with Dr. Dimitri J. Mavriplis (ICASE).

PO-SHU CHEN

Parallel Solution of Coupled Aeroelastic Problems

The accurate prediction of aeroelastic response is essential in the design of high performance aircraft. It requires solving the coupled fluid and structure equations simultaneously. The advance of the powerful parallel processor provides the ideal tool for this problem. The objectives of this research are to investigate a variety of different approaches for solving aeroelastic problems, to establish a proper module between structure and fluid simulations, to solve the aeroelastic response, and to research a better integration algorithm for communication between fluid and structure equations.

In the last several months, we have validated load transfer algorithms and the inverse isoparametric method by solving the HSCT model with the method developed by Prof. Farhat (University of Colorado). Currently, the MDO branch is using FASIT code developed by the Air Force/Georgia Institute of Technology. This code is compatible with most of the commercial software like NASTRAN, PLOT3D and CFL3D. However, the code is limited to structured fluid mesh only, and also has a poor user interface and uses outdated programming methodology. On the other hand, the Colorado package lacks the mechanism for integrating with these well-known software.

We would like to develop a new package to replace the function of FASIT. This new package will not only retain the compatibility with the popular software, but also use the latest integration algorithm in

college. There are two parts in this new package. The first part is called *matcher*, which will be used for searching the interface between structure and fluid mesh and handling the unstructured mesh. The second part handles the integration. Since *matcher* is an existing code module and parts of the integration routines have been developed and tested already, we will be able to deliver the package in a short time period.

This work is being done in collaboration with Tom Zang and Anthony Giunta (NASA Langley) and Professor Charbel Farhat (University of Colorado).

THOMAS W. CROCKETT

Application of Parallel and Distributed Computing to Visualization and Data Assimilation Problems in the Atmospheric Sciences

To implement the Vice President's vision of a Digital Earth, vast quantities of data from disparate sources must be integrated into an intuitive, accessible representation. NASA's Earth Science Enterprise sees Digital Earth as a promising framework for making much of its remote sensing data available to the scientific community and the general public. To implement the Digital Earth concept, many technologies will need to be brought to bear, among them visualization, networking, and high-performance computing.

We are exploring the potential for parallel and distributed computing and visualization techniques to contribute to the data processing and data assimilation requirements of Digital Earth. In particular, we are interested in the utility of parallel rendering methods to alleviate bottlenecks in the visualization process, to reduce the end-to-end time required for data processing and data assimilation, and/or to enable the production of new visualization products. Our initial focus is on atmospheric lidar data generated by Langley's spaceborne LITE experiment. We began with a series of experiments to characterize networking performance among the principle computer systems which will be involved in this project: ICASE's UltraSPARC workstations, the HPCCP-sponsored Origin2000 operated by ISSD, and the large Onyx2 visualization system in the Data Visualization and Animation Lab. Based on these results, we were able to optimize the networking code in our PGL rendering system, improving frame rates by about 80% (to 9 fps between ISSD and ICASE for a minimally-compressible, $512 \times 512 \times 24$ -bit animation sequence).

We are currently developing a PGL-based visualization application which combines a medium-resolution (9 km) elevation model of the Earth with a true-color surface map. Once the Earth model is completed, we will use it to benchmark rendering performance on the Origin2000 and compare it with DVAL's Onyx system. We will also add capabilities for importing and visualizing LITE data, referencing it to the space shuttle's track over the Earth. The ultimate goal is to develop a responsive, user-friendly system which will combine atmospheric data from a variety of sources to obtain a better understanding of the physical processes involved in the Earth's atmosphere.

Cross-Platform Portability and Performance of Parallel Rendering Algorithms

As parallel computer architectures continue to evolve in several distinct directions, the problem of developing portable parallel software becomes more urgent. Currently, message passing is the only widely supported programming paradigm which can span the three major architectures of interest: loosely-coupled clusters of PC's and workstations, tightly-coupled distributed-memory systems, and the emerging class of distributed-shared-memory (DSM) machines. An important question is whether the message passing approach can deliver satisfactory performance across this diverse set of computing platforms.

Over the last several months, we have studied this issue in detail within the context of parallel polygon rendering, using our PGL system as the driving application. Although PGL is undergoing continuous development, the basic structure of the communication algorithms has changed little since 1992. Hence we have roughly comparable performance measurements across four generations of parallel architectures, beginning with the Intel iPSC/860 and including the Intel Paragon, IBM SP2, Cray T3E, SGI Origin2000, and HP Exemplar. The first four are pure distributed-memory systems while the last two are distributed-shared-memory machines, with hardware support for direct addressing of remote memory locations. Although the communication architectures and processing power of the four distributed-memory systems are vastly different, our experiments show that they exhibit similar scalability, with parallel efficiencies in the 50-70% range for 128 processors on a standard benchmark scene. In contrast, the two DSM systems performed much less efficiently in large configurations, achieving parallel efficiencies of only 10-20% with 128 processors. A number of factors appear to contribute to these poor results, including contention for shared data structures, inefficient MPI implementations, the absence of message co-processors, and memory management and process scheduling issues.

These results suggest that message passing may not be an effective programming paradigm for communication-intensive asynchronous algorithms on DSM architectures. We are currently trying to determine to what extent performance can be improved within the context of the existing algorithms. Preliminary results from the HP Exemplar indicate that a factor of two may be possible by reducing the frequency of polling operations. However, this still falls short of the efficiencies achievable on distributed-memory systems. In the longer term, we plan to explore alternate shared memory formulations to determine whether the current problems are due solely to message passing overheads, or whether the DSM architectures themselves have inherent limitations for these kinds of problems.

STEPHEN GUATTERY

Pseudoinverses of Symmetric Matrices and Generalized Graph Embeddings

Connections between Laplacian spectra (specifically the second smallest eigenvalue λ_2 and its corresponding eigenvector, or, in the case of a zero Dirichlet boundary, the smallest eigenvalue λ_1 and its eigenvector) and properties of the corresponding graphs have applications in algorithms. A number of graph embedding techniques that give lower bounds on the smallest nontrivial eigenvalues of Laplacians exist (an example is the path resistance method recently introduced by Steve Guattery (ICASE), Tom Leighton (MIT), and Gary Miller (CMU), which works for both the graph (positive semidefinite) and zero boundary (positive definite) cases). A common attribute of embedding techniques is that they do not provide tight lower bounds; Guattery and Miller have recently shown that the gap in these bounds is a result of the representation of the problem. They showed that, by slightly modifying the representation, it is possible to construct a specific embedding, the current flow embedding, such that the matrix representation of the embedding can be used to construct a matrix whose eigenvalues and eigenvectors have an exact relationship to those of the Laplacian. In the Dirichlet boundary case, they showed that the embedding matrix can be used to construct the inverse of the Laplacian. These results have an interpretation in terms of resistive circuits and Kirchoff's and Ohm's laws. This result was generalized to show that, for any real symmetric matrix, its pseudoinverse can be constructed in terms of a generalization of the current flow embedding. The resistive circuit analogy is lost in this case, though the embedding still obeys generalizations of the electrical laws.

Most recently, this work was extended to cover all Hermitian matrices. An ICASE report has been submitted.

Current work is focusing on whether approximations to the generalized current flow embedding can be used to construct useful preconditioners for such systems continues.

DAVID E. KEYES

Parallel Implicit Solvers for Simulation of Multiscale Phenomena

The development and application of parallel implicit solvers for multiscale phenomena governed by PDEs are the chief objectives of our work. Newton-Krylov-Schwarz (NKS) methods have proven to be broadly applicable, architecturally versatile, and tunable for high performance on today's high-end commercial parallel platforms (e.g., Cray T3E, SGI Origin, IBM SP). Both structured-grid and unstructured-grid CFD legacy codes have been ported to such platforms and reasonable objectives for algorithmic convergence rate, parallel efficiency, and raw floating point performance have been met. However, architectural challenges have increased on the next generation of high-end machines, as represented, for instance, by the ASCI "blue" machines at Lawrence Livermore and Los Alamos National Laboratories, and also on the increasingly pervasive low-end parallel clusters consisting of commodity processors running Linux, known as "Beowulf"s. There is simultaneously a mandate to make parallel computing accessible to a wider group of computational engineers and physicists with complex applications, for whom preparation of a Jacobian linearization for Newton's method is impractical or intimidating. Our recent efforts have therefore concentrated on algorithmic adaptations of NKS methodology appropriate for the emerging architectures and on evaluation of new software tools and methodology to assist in future ports.

We categorize our approaches according to the obstacles presented by new computational environments and applications.

The ASCI "blue" machines are hybrid, distributed shared-memory (DSM) multiprocessors. One may ignore the local shared-memory capabilities at modest performance cost, and continue to treat each processor as the exclusive owner of a share of the overall memory, and this is the approach we have taken thus far. However, we are also now studying the OpenMP standard and other compiler-controlled means of multi-threading.

A chief concern at fine granularity parallelism or in the context of weakly connected clusters is processor idle time because of data starvation due ultimately to load imbalance. More generally, floating point work may be balanced, but progress towards the overall convergence of the computation is still imbalanced due to linear or nonlinear "stiffness." For these reasons, we continue to study the role of semi-synchronous Schwarz iteration outside of subdomain-scale Newton iterations, at least in early stages of the simulation. This leads to consideration of user-controlled task-level multi-threaded approaches, perhaps in conjunction with the compiler-controlled instruction-level form of the previous paragraph.

Apart from parallel scalability, per-node floating point performance has been a source of major consternation for users (and purchasers) of high-end machines. We have shown that attention to cache line reuse in the organization and ordering of grid-based data that is iteratively dragged up and down the memory system in a typical PDE code can make an order of magnitude difference in execution time, apart from parallelism, and have commenced an experimental program to study this effect via hardware event counters. Our ultimate aim is to apply formal optimization techniques to the layout of program data for optimal register and cache residency. We note that changes in the algorithm, itself, could make this task easier: for instance, high-order methods tend to produce local dense blocks with better concentration of flops within smaller working sets.

With a goal of lowering the barrier of moving legacy codes into the parallel implicit environment, we have been testing the automatic differentiation software packages from Argonne National Laboratory, ADIFOR and ADIC, and have found them to be relatively friendly on simple applications. However, we believe that the user will want to base Jacobian approximations on a discretization that is perhaps distinct from the high-fidelity discretization whose residual is ultimately being driven to zero by Newton's method, so such tools will not necessarily permit full automation.

We have also participated in the parallel porting of an unstructured-grid code based on the discontinuous Galerkin discretization, which was implemented in C++, via the object-oriented concept of inheritance of parallel attributes.

We will continue to pioneer nonlinear Schwarz methods in implicit parallel CFD, including developing metrics for nonlinearity and developing a software environment appropriate for less synchronous computation. We will work on improved static and dynamic partitioning methods, in order to load balance the sum of computation and communication work, rather than just the computation work, as at present. We will continue to work with the CFD user community and begin to work with the radiation transport user community to make Beowulf-style and HPCCP testbed use "friendly."

This work was done in collaboration with W. Kyle Anderson and Harold Atkins (NASA Langley), Dinesh Kaushik and Nilan Karunaratne (Old Dominion University), Abdelkader Baggag (Purdue University), William D. Gropp, Lois C. McInnes, and Barry F. Smith (Argonne National Laboratory).

GERALD LÜTTGEN

Equational Reasoning for User-defined Congruences in PVS

Proving that a system implementation meets its specification, both of which are given terms of an algebra, can be done by replacing subterms by equal terms, where the underlying notion of equality is described by a set of axioms which are sound for the congruence of interest. This kind of equational reasoning or rewriting is popular among engineers and is also the core of process algebras, which allow one to reason about the communication behavior of concurrent systems. In order to support equational reasoning, theorem provers have been used to perform elementary rewriting steps and to apply proof strategies. However, SRI's theorem prover PVS, which is used at NASA LaRC, does not (explicitly) support equational reasoning for user-defined congruences. Our objective is to show how this kind of reasoning can be implemented implicitly in PVS, thereby invalidating conjectures of other researchers.

The well-known process algebra CCS, together with its behavioral congruence bisimulation \sim , serves as an example. To implement CCS in PVS we have: (i) defined process terms as a PVS datatype, (ii) encoded operational rules in a single function, (iii) defined bisimulation using PVS's subtype mechanism, and (iv) formalized CCS axioms as lemmas. Now elementary rewriting steps can be automated as follows. Let $C[P] \sim Q$ be the proof goal, where $C[P]$ denotes a process term containing P , and let $P \sim P'$ be an axiom. For substituting P by P' in $C[P] \sim Q$ we refer to the transitivity of \sim , i.e. $(\exists R. C[P] \sim R \wedge R \sim Q) \implies C[P] \sim Q$. By instantiating R with $C[P']$ and using compositionality to obtain $C[P] \sim C[P']$, one infers $C[P'] \sim Q$, as desired. These proof steps can be fully captured as a PVS proof strategy. Since our approach is generic and, in contrast to related work, does not sacrifice any given algebraic structure, we have achieved our goal.

Now that PVS is able to perform elementary rewriting steps for user-defined congruences, the question arises whether equational reasoning can be further automated. This work was done in collaboration with Paul S. Miner (NASA Langley).

KWAN-LIU MA

Parallel Rendering of 3D AMR Data

Adaptive Mesh Refinement (AMR) is a powerful tool for modeling many important scientific and engineering problems. However, visualization tools for three-dimensional AMR data are not generally available. The adaptive nature of the embedded mesh demands sophisticated visualization calculations. To obtain high accuracy visualization results, rendering three-dimensional AMR data on an average graphics workstation can be painfully slow. A desirable approach is thus to distribute both the data and rendering tasks to multiple computers. Since massively parallel computers are becoming more and more accessible, the goal of this research is to develop an efficient three-dimensional visualization strategy for AMR data by utilizing the state-of-the-art massively parallel computers such as the SGI/Cray T3E and SGI Origin 2000.

In this work, we have mainly considered the type of data generated from applications using the PARAMESH package developed at NASA's Goddard Space Flight Center. With PARAMESH, a hierarchy of sub-grids can be generated to cover the computational domain, with spatial resolution varying to satisfy the demands of the application. These sub-blocks form the nodes of a tree data-structure (quadtree in 2D or octree in 3D). Each grid block has a logically Cartesian mesh, and the index ranges are the same for every block. We could make use of an existing parallel renderer by uniformly resampling the data before rendering, but then both the storage and rendering overhead becomes very high. A better approach is to render the AMR data directly; that is, sampling is adapted to the resolution of each block. However, because solution data are cell-centered, we must first convert the raw data to being vertex-centered for facilitating the volume rendering calculations.

We have completed the data conversion program in which care must be taken for interpolation between blocks. We have also completed a prototype of the parallel renderer using MPI. We are currently in the process of collecting test data. Preliminary results demonstrate very good rendering rates, seconds to a few minutes when using a cluster of eight SUN workstations, and subseconds to seconds when using the Cray T3E for datasets which, if uniformly sampled, one would consist of 512^3 voxels and the other 1024^3 voxels.

As we see an increasing use of the PARAMESH package in the modeling of many nationally-relevant scientific problems, the strategies we produce in this research will benefit users of PARAMESH (and other AMR technologies) in a significant way.

Future work includes creating a graphics user interface for the renderer and other preprocess steps, and developing the capability of viewing the computational grid itself and displaying solution data superposed on a highly-magnified grid which is of very strong interest to application developers for tuning and debugging purposes.

PIYUSH MEHROTRA

Arcade: A Distributed Computing Environment for ICASE

Distributed heterogeneous computing is being increasingly applied to a variety of large size computational problems. Such computations, for example, the multidisciplinary design optimization of an aircraft, generally consist of multiple heterogeneous modules interacting with each other to solve the problem at hand. Such applications are generally developed by a team in which each discipline is the responsibility of experts in the field. The objective of this project is to develop a GUI-based environment which supports the multi-user design of such applications and their execution and monitoring in a heterogeneous environment consisting of a network of workstations, specialized machines, and parallel architectures.

We have been implementing a Java-based three-tier prototype system which supports a thin client interface for the design and execution of multi-module codes. The middle tier consists of logic to process the user input and also to manage the resource controllers which comprise the third tier. In the last few months we have concentrated on providing support for cross-domain execution of sub-modules of the application. That is, the user can specify that a particular module be executed on a resource which is in a different domain. In such a situation, the execution controller has to transfer the data to the remote resource before requesting that execution be started. We designed several different mechanisms for the interaction between the domains based on straight TCP/IP and also on HTTP. We are currently examining these different approaches both from the point of view of performance and security.

We have also been studying the use of latest technology developed by Sun called JINI for resource monitoring and management. JINI allows independent resources to announce their presence and current status to a central server. We are currently in the process of building a client interface to allow users to monitor the status of the various resources.

In the future we will be adding other features to the system including the interactive specification of execution resources. We are also planning to examine the use of Globus, a distributed program execution toolkit, for supporting remote execution of modules.

This work is being done in collaboration with K. Maly, A. Al-theneyan, M. Zubair (Old Dominion University).

Multithreaded System for Distributed Environments

Traditionally, lightweight threads are supported only within the single address space of a process, or in shared memory environments with multiple processes. Likewise, interprocess communication systems do not currently allow messages to be sent directly to entities within a process. The objective of this project is build a system which combines standard interfaces for lightweight threads, pthreads, and interprocess communication, MPI, to support point-to-point communication between any two threads in a distributed memory system.

The Chant runtime system has been built using layers: point-to-point communication, remote service requests, remote thread operations. In the last year we have added a layer to support load balancing via migration of threads. In contrast to other thread migration systems, we provide migration in the presence of pointers, i.e., along with the stack, the thread heap is also migrated. This allows the pointers to point to valid data even after the migration. The load balancing layer provides facilities for computing the current workloads of the processors, figuring which threads should be moved and to what processors. The system also keeps track of the amount of communication being generated by each thread so that it can be used in the decision making process. The underlying thread migration layer takes care of the actual motion of the threads. The system is designed such that default routines can be replaced by the user with specialized routines for carrying out the various functions.

In the last few months, we have been optimizing the implementation based on our tests using several applications. These include a simple adaptive quadrature code a traveling salesman code and an image rendering code. In each of these cases we have shown that there is an overall gain in performance due to thread-based load balancing. That is, the overhead due to multi-threading, thread migration and load management, as implemented in our system are at acceptable levels.

This work was done in collaboration with D. Cronk (Ph.D. student, The College of William & Mary).

Integrating High Performance Fortran and OpenMP

Recently a proposal was put forth for a set of language extensions to Fortran and C based upon a fork-join model of parallel execution; called **OpenMP**, it aims to provide a portable shared memory programming interface for shared memory and low latency systems. However, these extensions ignore the issue of data locality which becomes a performance issue on shared address space machines which use a physically distributed memory system. We have been investigating how the two models, HPF and OpenMP can be used together to write programs which exploit the full capabilities of such systems.

We have examined several approaches to this problem. One approach is to use the HPF extrinsic procedures to call OpenMP routines. This allows both models to be separately used. A more integrated approach is provide HPF-like mapping directives in OpenMP so as to be able to control the data locality. We have examined the OpenMP constructs and have suggested how they can be extended so as to allow directives specifying the distribution of data and the work associated with the data can be explicitly specified by the user. In particular, we have defined a hybrid model of parallel computation by adding distribution directives, generalizing OpenMP's privatization concept and introducing the independent loop.

We are planning to implement these constructs in order to test their performance and efficacy.

This work is being done in collaboration with K. Roe (Ph.D. student, The College of William & Mary), B. Chapman (Southampton University) and H. Zima (University of Vienna).

CAN ÖZTURAN

Distributed Environment and Algorithms for Unstructured Meshes

The aim of this project is to develop a software environment and algorithms for unstructured meshes. Dynamic unstructured meshes which adaptively evolve during runtime require complicated data structures and algorithms for data management, mesh refinement/coarsening, and load balancing.

Development and testing work for Parallel Mesh Environment (PME) was continued. In particular, previous implementation for the Bansch/Joe algorithm for tetrahedral mesh refinement was debugged and tested. A POSIX multi-threaded implementation of the data-parallel refinement algorithm given in previous work was programmed. Extensive collaborative work was carried out with Abdelkader Baggag (Purdue University) to document and improve the parallel aeroacoustics solver.

PME is currently implemented in C. Although, much effort was spent in maintaining an object-oriented-like design and interface, a C++ implementation will simplify coding and offer several advantages. Therefore, a complete revision of PME implementation is planned. Additionally, simple data refinement algorithms for tetrahedral meshes similar to previous work is under investigation.

ALEX POTHEN

Enhancing the Cache Performance of Irregular Computations

Modern computer architectures improve performance by executing a number of instructions in parallel: they achieve this by pipelining instructions, replicating functional units, issuing multiple instructions per clock cycle, and predicting the direction of branches. However, since execution times of instructions improve at a faster rate than memory access times, feeding the CPU with data operands is the bottleneck in many irregular computations. This bottleneck is alleviated by means of caches, small fast memories to which data can be moved from the larger and slower primary memory, and from which data can be moved to the

CPU. We study how irregular computations (e.g., unstructured mesh or sparse matrix algorithms) can be organized so that most of the data accesses can be performed within the cache for high performance.

We have studied the cache performance of the kernel of an unstructured mesh code for solving Euler's equations using an "on-the-fly" cache simulation tool. This tool, called FastCache, was developed by Alvin Lebeck and David Wood (University of Wisconsin). The cache misses are divided into three classes: compulsory (first reference), capacity, and conflict. We have reordered the data accesses in this computation with the Cuthill-McKee, Sloan, and Nested Dissection algorithms. The cache performance of irregular codes is greatly enhanced by these orderings, and factors of two to three reduction in run-time are obtained by reordering the cache accesses. We have looked at the number of cache misses generated by each statement in the kernel to understand the reasons why the reordering algorithms improve cache performance. We conclude that the CM and Sloan reorderings improve spatial locality, while Nested Dissection improves temporal locality. We have also shown that as problem sizes increase, any reordering algorithm must partition the problem to maintain good cache performance. Our findings lead to general principles for improving the cache performance of irregular computations.

This study is in its early stages. We are examining how orderings based on space-filling curves enhance cache performance, and are attempting to develop analytical models for the cache performance of scientific computing algorithms. The performance of irregular codes on multilevel memory hierarchies is another new direction to explore.

This work was done in collaboration with Shengnian Ye (a student at Old Dominion University) and Dimitri Mavriplis (ICASE).

Object-oriented Design of Sparse Solvers

Direct methods for solving systems of linear equations employ sophisticated combinatorial and algebraic algorithms that contribute to software complexity, and hence it is natural to consider object-oriented design (OOD) in this context. We have continued to create software for solving sparse systems of linear equations by direct methods employing OOD.

OOD manages complexity by means of decomposition and abstraction. We decompose our software into two main types of objects: structural objects corresponding to data structures, and algorithmic objects corresponding to algorithms. This design decouples data structures from algorithms, permitting a user to experiment with different algorithms and different data structures, and if necessary develop new algorithms and data structures. We have implemented a family of minimum degree ordering algorithms using this design paradigm, and thus provided a laboratory for users to experiment with recent enhancements and algorithmic variants. We have also implemented a direct solver for symmetric positive definite and indefinite problems. The increased flexibility in OOD could come at the cost of some loss in efficiency. We have made careful trade-offs in our software to achieve the benefits of OOD without sacrificing efficiency. The running times of our C++ code for the symmetric positive definite solver compare quite favorably with existing Fortran 77 codes. We have reported our results at the SCITOOLS98 (Oslo) and ISCOPE98 (Santa Fe) conferences this year. This is the first object-oriented implementation of sparse direct methods of which we know.

We are continuing to create parallel and out-of-core solvers using OOD. Our current serial code is available with an interface to the PETSc toolkit from Argonne National Labs, and as a stand-alone code.

This work was done in collaboration with Florin Dobrian and Gary Kurfert (Ph.D. students at Old Dominion University).

Fast computation of robust preconditioners is a priority for solving large systems of equations on unstructured grids and in other applications. We have developed new algorithms and software that can compute incomplete factorization preconditioners for high level fill in time proportional to the number of floating point operations and memory accesses.

We have developed a structure theory based on paths in the adjacency graph of the matrix to predict where zero elements become nonzeros in incomplete factorization (fill elements). A level function is used in incomplete factorization to control the number of fill elements, and we relate the level of fill to lengths of appropriately defined paths in the adjacency graph. This result permits us to search in the neighborhood of a vertex in the graph to predict all fill elements associated with that vertex. We have designed two variants of these algorithms and have proved that they have a smaller running time complexity than currently used algorithms for computing incomplete factorizations. The more efficient algorithms make use of the concept of transitive reduction of directed graphs (symmetric problems) and symmetric reduction of directed graphs (unsymmetric problems) in order to search for paths in smaller graphs. Our implementation in C shows that the new algorithms are faster than implementations available earlier.

We are investigating parallel implementations of incomplete factorizations. Our serial code is available with an interface to the PETSc toolkit from Argonne; it adds functionality in the area of preconditioners to PETSc.

This work was done in collaboration with David Hysom (Ph.D. student at Old Dominion University and ICASE).

HYEON-JU YOON

Parallel Isosurface Extraction

Isosurface extraction is one of the essential techniques which can help to understand and visualize volume data. The more processing power grows, the more data are generated. Parallelized implementation is a good approach for reducing the computation time for isosurface extraction from large volume of data.

As a basis for parallelization, we chose the ‘marching cube algorithm,’ and find an appropriate optimization method to parallelize. Several optimization methods which use reorganization of original data, have been proposed to avoid unnecessary or redundant calculations; for example, octree, span space, and interval tree. These methods have shown great performance gains with rather small storage overhead when compared to the original marching cube algorithm. But, their data structures are usually spatially unbalanced, so they often cause load imbalance and result in performance degradation when they are executed on parallel machines. We also have to consider the time varying data set in the real time environment. Preprocessing time for reorganization of volume data is not a big overhead if the additional data structure can be reused several times, but it may be unnecessary for the time varying data. Currently, we have implemented sequential versions of these methods and have also collected statistics for analysis.

Based upon these sequential versions, we plan to apply these methods to construct the parallel isosurface extraction library. Additionally, we will also study new techniques for optimization and unstructured data.

HANS ZIMA

High-level Support for Semi-structured Algorithms

High Performance Fortran (HPF) provides a high-level language interface for programming scalable parallel architectures. While the usefulness of the language for regular problems has been generally accepted, its efficient applicability to irregular and semi-regular problems is still a research issue. The objective of this project is to study the efficacy of the HPF directives for specifying the parallelism in codes using semi-structured data structures.

This project is based upon the de-facto standard HPF-2 and its enhancement in form of the HPF+ language which was developed in an ESPRIT project based on joint work by ICASE and the University of Vienna. In particular, we have studied the representation of various types of grid collections in HPF. We have also analyzed the degree of parallelism that can be obtained for multiblock, semi-coarsening multigrid, and structured AMR algorithms under a range of data distribution strategies.

We plan to implement a set of kernel algorithms using the Vienna Fortran Compiler (VFC) to compare the performance of the different distribution strategies.

This work was done in collaboration with Piyush Mehrotra (ICASE).

Implementing the Opus Language

Multidisciplinary applications require an integration of task parallelism with data parallelism. In previous work, ICASE and the University of Vienna have jointly developed the language OPUS, which provides an object-based layer for the coordination of multiple HPF programs operating in a parallel or heterogeneous environment. The objective of this project is to design and implement a runtime system for Opus.

During this reporting period, a detailed implementation design was made and a prototype implementation in the framework of distributed-memory machines and workstation clusters was completed.

Future work will extend the present design and implementation to take into account the requirements of heterogeneous distributed systems. Furthermore, the interface between Opus and the Java language will be examined.

This work was done in collaboration with Piyush Mehrotra (ICASE) and Erwin Laure (University of Vienna).

REPORTS AND ABSTRACTS

Canuto, Claudio, Anita Tabacco, and Karsten Urban: *The wavelet element method part II: Realization and additional features in 2D and 3D*. ICASE Report No. 98-17, (NASA/CR-1998-207637), April 3, 1998, 42 pages. Submitted to Applied and Computational Harmonic Analysis.

The Wavelet Element Method (WEM) provides a construction of multiresolution systems and biorthogonal wavelets on fairly general domains. These are split into subdomains that are mapped to a single reference hypercube. Tensor products of scaling functions and wavelets defined on the unit interval are used on the reference domain. By introducing appropriate matching conditions across the interelement boundaries, a globally continuous biorthogonal wavelet basis on the general domain is obtained. This construction does not uniquely define the basis functions but rather leaves some freedom for fulfilling additional features.

In this paper we detail the general construction principle of the WEM to the 1D, 2D and 3D cases. We address additional features such as symmetry, vanishing moments and minimal support of the wavelet functions in each particular dimension. The construction is illustrated by using biorthogonal spline wavelets on the interval.

Loncaric, Josip: *Optimal control of unsteady Stokes flow around a cylinder and the sensor/actuator placement problem*. ICASE Report No. 98-18, (NASA/CR-1998-207680), June 4, 1998, 23 pages. To appear in the Proceedings of the AFOSR Workshop on Optimal Design and Control.

Effective placement of sensors and actuators is of crucial importance in flow control. Instead of using combinatorial search to identify optimal locations, we pose a related problem of polynomial complexity. If one could sense everything and actuate everywhere, what should one do? Using the unsteady 2D Stokes flow around a cylinder as an example, we obtain the analytic solution of an optimal distributed control problem and describe its spatial structure. At low circumferential wavenumbers or close to the cylinder wall, boundary vortex generators are shown to be more effective than colocated vorticity damping. This analytic solution has also been used to test numerical methods, demonstrating the importance of using discretization which resolves all eigenfunctions of interest.

Nordstrom, Jan, and Mark H. Carpenter: *Boundary and interface conditions for high order finite difference methods applied to the Euler and Navier-Stokes equations*. ICASE Report No. 98-19, (NASA/CR-1998-207681), June 4, 1998, 27 pages. Submitted to the Journal of Computational Physics.

Boundary and interface conditions for high order finite difference methods applied to the constant coefficient Euler and Navier-Stokes equations are derived. The boundary conditions lead to strict and strong stability. The interface conditions are stable and conservative even if the finite difference operators and mesh sizes vary from domain to domain. Numerical experiments show that the new conditions also lead to good results for the corresponding nonlinear problems.

Mavriplis, Dimitri J.: *Three-dimensional high-lift analysis using a parallel unstructured multigrid solver*. ICASE Report No. 98-20, (NASA/CR-1998-207682), June 17, 1998, 20 pages.

A directional implicit unstructured agglomeration multigrid solver is ported to shared and distributed memory massively parallel machines using the explicit domain-decomposition and message-passing approach.

Because the algorithm operates on local implicit lines in the unstructured mesh, special care is required in partitioning the problem for parallel computing. A weighted partitioning strategy is described which avoids breaking the implicit lines across processor boundaries, while incurring minimal additional communication overhead. Good scalability is demonstrated on a 128 processor SGI Origin 2000 machine and on a 512 processor CRAY T3E machine for reasonably fine grids. The feasibility of performing large-scale unstructured grid calculations with the parallel multigrid algorithm is demonstrated by computing the flow over a partial-span flap wing high-lift geometry on a highly resolved grid of 13.5 million points in approximately 4 hours of wall clock time on the CRAY T3E.

Adumitroaie, V., J.R. Ristorcelli, and D.B. Taulbee: *Progress in Favré-Reynolds stress closures for compressible flows*. ICASE Report No. 98-21, (NASA/CR-1998-208423), June 8, 1998, 42 pages. Submitted to Physics of Fluids.

A closure for the compressible portion of the pressure-strain covariance is developed. It is shown that, within the context of a pressure-strain closure assumption linear in the Reynolds stresses, an expression for the pressure-dilatation can be used to construct a representation for the pressure-strain. Additional closures for the unclosed terms in the Favré-Reynolds stress equations involving the mean acceleration are also constructed. The closures accommodate compressibility corrections depending on the magnitude of the turbulent Mach number, the mean density gradient, the mean pressure gradient, the mean dilatation, and, of course, the mean velocity gradients. The effects of the compressibility corrections are consistent with current DNS results. Using the compressible pressure-strain and mean acceleration closures in the Favré-Reynolds stress equations an algebraic closure for the Favré-Reynolds stresses is constructed. Noteworthy is the fact that, *in the absence of mean velocity gradients*, the mean density gradient produces Favré-Reynolds stresses in accelerating mean flows. Computations of the mixing layer using the compressible closures developed are described. Full Reynolds stress closure and two-equation algebraic models are compared to laboratory data. The mixing layer configuration computations are compared to laboratory data; since the laboratory data for the turbulence stresses is inconsistent, this comparison is inconclusive. Comparisons for the spread rate reduction indicate a sizable decrease in the mixing layer growth rate.

Ma, Kwan-Liu, Diann Smith, Ming-Yun Shih, and Han-Wei Shen: *Efficient encoding and rendering of time-varying volume data*. ICASE Report No. 98-22, (NASA/CR-1998-208424), June 8, 1998, 18 pages. Submitted to Visualization '98 Conference.

Visualization of time-varying volumetric data sets, which may be obtained from numerical simulations or sensing instruments, provides scientists insights into the detailed dynamics of the phenomenon under study. This paper describes a coherent solution based on quantization, coupled with octree and difference encoding for visualizing time-varying volumetric data. Quantization is used to attain voxel-level compression and may have a significant influence on the performance of the subsequent encoding and visualization steps. Octree encoding is used for spatial domain compression, and difference encoding for temporal domain compression. In essence, neighboring voxels may be fused into macro voxels if they have similar values, and subtrees at consecutive time steps may be merged if they are identical.

The software rendering process is tailored according to the tree structure and the volume visualization process. With the tree representation, selective rendering may be performed very efficiently. Additionally, the I/O costs are reduced. With these combined savings, a higher level of user interactivity is achieved. We

have studied a variety of time-varying volume datasets, performed encoding based on data statistics, and optimized the rendering calculations wherever possible. Preliminary tests on workstations have shown in many cases tremendous reduction by as high as 90% in both storage space and inter-frame delay.

Guattery, Stephen, and Gary L. Miller: *Graph embeddings and Laplacian eigenvalues*. ICASE Report No. 98-23, (NASA/CR-1998-208425), July 9, 1998, 20 pages. Submitted to SIAM Journal on Matrix Analysis and Applications.

Graph embeddings are useful in bounding the smallest nontrivial eigenvalues of Laplacian matrices from below. For an $n \times n$ Laplacian, these embedding methods can be characterized as follows: The lower bound is based on a clique embedding into the underlying graph of the Laplacian. An embedding can be represented by a matrix Γ ; the best possible bound based on this embedding is $n/\lambda_{\max}(\Gamma^T \Gamma)$. However, the best bounds produced by embedding techniques are not tight; they can be off by a factor proportional to $\log^2 n$ for some Laplacians.

We show that this gap is a result of the representation of the embedding: by including edge directions in the embedding matrix representation Γ , it is possible to find an embedding such that $\Gamma^T \Gamma$ has eigenvalues that can be put into a one-to-one correspondence with the eigenvalues of the Laplacian. Specifically, if λ is a nonzero eigenvalue of either matrix, then n/λ is an eigenvalue of the other. Simple transformations map the corresponding eigenvectors to each other. The embedding that produces these correspondences has a simple description in electrical terms if the underlying graph of the Laplacian is viewed as a resistive circuit. We also show that a similar technique works for star embeddings when the Laplacian has a zero Dirichlet boundary condition, though the related eigenvalues in this case are reciprocals of each other. In the Dirichlet boundary case, the embedding matrix Γ can be used to construct the inverse of the Laplacian. Finally, we connect our results with previous techniques for producing bounds, and provide an illustrative example.

Gropp, W.D., D.E. Keyes, L.C. McInnes, and M.D. Tidriri: *Globalized Newton-Krylov-Schwarz algorithms and software for parallel implicit CFD*. ICASE Report No. 98-24, (NASA/CR-1998-208435), August 20, 1998, 42 pages. To be submitted to the International Journal of Supercomputer Applications and High Performance Computing.

Implicit solution methods are important in applications modeled by PDEs with disparate temporal and spatial scales. Because such applications require high resolution with reasonable turnaround, "routine" parallelization is essential. The pseudo-transient matrix-free Newton-Krylov-Schwarz (Psi-NKS) algorithmic framework is presented as an answer. We show that, for the classical problem of three-dimensional transonic Euler flow about an M6 wing, Psi-NKS can simultaneously deliver: globalized, asymptotically rapid convergence through adaptive pseudo-transient continuation and Newton's method; reasonable parallelizability for an implicit method through deferred synchronization and favorable communication-to-computation scaling in the Krylov linear solver; and high per-processor performance through attention to distributed memory and cache locality, especially through the Schwarz preconditioner. Two discouraging features of Psi-NKS methods are their sensitivity to the coding of the underlying PDE discretization and the large number of parameters that must be selected to govern convergence. We therefore distill several recommendations from our experience and from our reading of the literature on various algorithmic components of Psi-NKS, and we describe a freely available, MPI-based portable parallel software implementation of the solver employed here.

Interrante, Victoria, and Chester Grosch: *Recent advances in visualizing 3D flow with LIC*. ICASE Report No. 98-26, (NASA/CR-1998-208437), July 29, 1998, 17 pages. To appear in IEEE Computer Graphics and Applications.

Line Integral Convolution (LIC), introduced by Cabral and Leedom in 1993, is an elegant and versatile technique for representing directional information via patterns of correlation in a texture. Although most commonly used to depict 2D flow, or flow over a surface in 3D, LIC methods can equivalently be used to portray 3D flow through a volume. However, the popularity of LIC as a device for illustrating 3D flow has historically been limited both by the computational expense of generating and rendering such a 3D texture and by the difficulties inherent in clearly and effectively conveying the directional information embodied in the volumetric output textures that are produced. In an earlier paper, we briefly discussed some of the factors that may underlie the perceptual difficulties that we can encounter with dense 3D displays and outlined several strategies for more effectively visualizing 3D flow with volume LIC. In this article, we review in more detail techniques for selectively emphasizing critical regions of interest in a flow and for facilitating the accurate perception of the 3D depth and orientation of overlapping streamlines, and we demonstrate new methods for efficiently incorporating an indication of orientation into a flow representation and for conveying additional information about related scalar quantities such as temperature or vorticity over a flow via subtle, continuous line width and color variations.

Zannetti, Luca, and Angelo Iollo: *On the circulation manifold for two adjacent large lifting sections*. ICASE Report No. 98-27, (NASA/CR-1998-208447), July 30, 1998, 21 pages. Submitted to ZAMM.

The circulation functional relative to two adjacent lifting sections is studied for two cases. In the first case we consider two adjacent circles. The circulation is computed as a function of the displacement of the secondary circle along the axis joining the two centers and of the angle of attack of the secondary circle. The gradient of such functional is computed by deriving a set of elliptic functions with respect both to their argument and to their period. In the second case studied, we considered a wing-flap configuration. The circulation is computed by some implicit mappings, whose differentials with respect to the variation of the geometrical configuration in the physical space are found by divided differences. Configurations giving rise to local maxima and minima in the circulation manifold are presented.

Mahalov, Alex, Basil Nicolaenko, and Ye Zhou: *Energy spectra of strongly stratified and rotating turbulence*. ICASE Report No. 98-28, (NASA/CR-1998-208448), July 30, 1998, 11 pages. To appear in Physical Review E.

Turbulence under strong stratification and rotation is usually characterized as quasi-two dimensional turbulence. We develop a “quasi-two dimensional” energy spectrum which changes smoothly between the Kolmogorov $-5/3$ law (no stratification), the -2 scalings of Zhou for the case of strong rotation, as well as the -2 scalings for the case of strong rotation and stratification. For strongly stratified turbulence, the model may give the -2 scaling predicted by Herring; and the $-5/3$ scaling indicated by some mesoscale observations.

Zhou, Ye, Charles G. Speziale, and Robert Rubinstein: *Some remarks concerning recent work on rotating turbulence*. ICASE Report No. 98-29, (NASA/CR-1998-208449), July 30, 1998, 10 pages. Submitted to Physics of Fluids.

A recent paper on rotating turbulence by Canuto and Dubovikov is examined from both an historical and scientific perspective. It is first shown that their claim of finding a new energy spectrum scaling is inaccurate; such a scaling law has been published in the literature by other authors using the same physical assumptions. Canuto and Dubovikov actually only offered a different estimate for the constant. Finally, it is demonstrated that the alternative model for the dissipation rate transport equation proposed by Canuto and Dubovikov does not have the desired physical features in rotating isotropic turbulence. It is physically inconsistent in both the weak and strong rotation limits.

Iollo, Angelo, and Manuel D. Salas: *On the propagation of small perturbations in two simple aeroelastic systems*. ICASE Report No. 98-30, (NASA/CR-1998-208457), July 30, 1998, 14 pages. Submitted to Journal of Sound and Vibration.

In this paper we investigate the wave propagation patterns for two simple flow-structure problems. We focus on the study of the propagation speeds of the waves in the fluid and in the structure, as the rigidity of the structure and the Mach number of the undisturbed flow are changing. Some implications concerning the sound emission by inhomogeneities eventually present in the structure are discussed.

Lewis, Robert Michael, and Virginia Torczon: *A globally convergent augmented Lagrangian pattern search algorithm for optimization with general constraints and simple bounds*. ICASE Report No. 98-31, (NASA/CR-1998-208458), August 20, 1998, 20 pages. Submitted to SIAM Journal on Optimization.

We give a pattern search adaptation of an augmented Lagrangian method due to Conn, Gould, and Toint. The algorithm proceeds by successive bound constrained minimization of an augmented Lagrangian. In the pattern search adaptation we solve this subproblem approximately using a bound constrained pattern search method. The stopping criterion proposed by Conn, Gould, and Toint for the solution of this subproblem requires explicit knowledge of derivatives. Such information is presumed absent in pattern search methods; however, we show how we can replace this with a stopping criterion based on the pattern size in a way that preserves the convergence properties of the original algorithm. In this way we proceed by successive, inexact, bound constrained minimization without knowing exactly how inexact the minimization is. So far as we know, this is the first provably convergent direct search method for general nonlinear programming.

Hu, Changqing, and Chi-Wang Shu: *Weighted essentially non-oscillatory schemes on triangular meshes*. ICASE Report No. 98-32, (NASA/CR-1998-208459), July 30, 1998, 34 pages. Submitted to the Journal of Computational Physics.

In this paper we construct high order weighted essentially non-oscillatory (WENO) schemes on two dimensional unstructured meshes (triangles) in the finite volume formulation. We present third order schemes using a combination of linear polynomials, and fourth order schemes using a combination of quadratic polynomials. Numerical examples are shown to demonstrate the accuracies and robustness of the methods for shock calculations.

Torczon, Virginia, and Michael W. Trosset: *Using approximations to accelerate engineering design optimization*. ICASE Report No. 98-33, (NASA/CR-1998-208460), September 2, 1998, 20 pages. To appear in the Proc. of the 7th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization.

Optimization problems that arise in engineering design are often characterized by several features that hinder the use of standard nonlinear optimization techniques. Foremost among these features is that the functions used to define the engineering optimization problem often are computationally intensive. Within a standard nonlinear optimization algorithm, the computational expense of evaluating the functions that define the problem would necessarily be incurred for each iteration of the optimization algorithm.

Faced with such prohibitive computational costs, an attractive alternative is to make use of surrogates within an optimization context since surrogates can be chosen or constructed so that they are typically much less expensive to compute. For the purposes of this paper, we will focus on the use of algebraic approximations as surrogates for the objective.

In this paper we introduce the use of so-called merit functions that explicitly recognize the desirability of improving the current approximation to the objective during the course of the optimization. We define and experiment with the use of merit functions chosen to simultaneously improve both the solution to the optimization problem (the objective) and the quality of the approximation. Our goal is to further improve the effectiveness of our general approach without sacrificing any of its rigor.

Guattery, Stephen: *Graph embeddings, symmetric real matrices, and generalized inverses*. ICASE Report No. 98-34, (NASA/CR-1998-208462), August 5, 1998, 16 pages.

Graph embedding techniques for bounding eigenvalues of associated matrices have a wide range of applications. The bounds produced by these techniques are not in general tight, however, and may be off by a $\log^2 n$ factor for some graphs. Guattery and Miller showed that, by adding edge directions to the graph representation, they could construct an embedding called the current flow embedding, which embeds each edge of the guest graph as an electric current flow in the host graph. They also showed how this embedding can be used to construct matrices whose nonzero eigenvalues had a one-to-one correspondence to the reciprocals of the eigenvalues of the generalized Laplacians. For the Laplacians of graphs with zero Dirichlet boundary conditions, they showed that the current flow embedding could be used generate the inverse of the matrix. In this paper, we generalize the definition of graph embeddings to cover all symmetric matrices, and we show a way of computing a generalized current flow embedding. We prove that, for any symmetric matrix A , the generalized current flow embedding of the orthogonal projector for the column space of A into A can be used to construct the generalized inverse, or pseudoinverse, of A . We also show how these results can be extended to cover Hermitian matrices.

Patten, James M., and Kwan-Liu Ma: *A graph based interface for representing volume visualization results*. ICASE Report No. 98-35, (NASA/CR-1998-208468), August 20, 1998, 17 pages. To appear in the Proceedings of Graphics Interface '98.

This paper discusses a graph based user interface for representing the results of the volume visualization process. As images are rendered, they are connected to other images in a graph based on their rendering parameters. The user can take advantage of the information in this graph to understand how certain rendering parameter changes affect a dataset, making the visualization process more efficient. Because the

graph contains more information than is contained in an unstructured history of images, the image graph is also helpful for collaborative visualization and animation.

Lewis, Robert Michael, Anthony T. Patera, and Jaume Peraire: *A posteriori finite element bounds for sensitivity derivatives of partial-differential-equation outputs*. ICASE Report No. 98-36, (NASA/CR-1998-208469), August 20, 1998, 23 pages. Submitted to Finite Elements in Analysis and Design.

We present a Neumann-subproblem a posteriori finite element procedure for the efficient and accurate calculation of rigorous, “constant-free” upper and lower bounds for sensitivity derivatives of functionals of the solutions of partial differential equations. The design motivation for sensitivity derivative error control is discussed; the a posteriori finite element procedure is described; the asymptotic bounding properties and computational complexity of the method are summarized; and illustrative numerical results are presented.

Llorente, Ignacio M., and N. Duane Melson: *Robust multigrid smoothers for three dimensional elliptic equations with strong anisotropies*. ICASE Report No. 98-37, (NASA/CR-1998-208700), August 24, 1998, 44 pages. To be submitted to the Conference Proceedings of “The Iterative Methods Conference.”

We discuss the behavior of several plane relaxation methods as multigrid smoothers for the solution of a discrete anisotropic elliptic model problem on cell-centered grids. The methods compared are plane Jacobi with damping, plane Jacobi with partial damping, plane Gauss-Seidel, plane zebra Gauss-Seidel, and line Gauss-Seidel. Based on numerical experiments and local mode analysis, we compare the smoothing factor of the different methods in the presence of strong anisotropies. A four-color Gauss-Seidel method is found to have the best numerical and architectural properties of the methods considered in the present work. Although alternating direction plane relaxation schemes are simpler and more robust than other approaches, they are not currently used in industrial and production codes because they require the solution of a two-dimensional problem for each plane in each direction. We verify the theoretical predictions of Thole and Trottenberg that an exact solution of each plane is not necessary and that a single two-dimensional multigrid cycle gives the same result as an exact solution, in much less execution time. Parallelization of the two-dimensional multigrid cycles, the kernel of the three-dimensional implicit solver, is also discussed. Alternating-plane smoothers are found to be highly efficient multigrid smoothers for anisotropic elliptic problems.

Qian, Yue-Hong, and Ye Zhou: *Complete Galilean-invariant lattice BGK models for the Navier-Stokes equation*. ICASE Report No. 98-38, (NASA/CR-1998-208701), August 28, 1998, 11 pages. To appear in Europhysics Letters.

Galilean invariance has been an important issue in lattice-based hydrodynamics models. Previous models concentrated on the nonlinear advection term. In this paper, we take into account the nonlinear response effect in a systematic way. Using the Chapman-Enskog expansion up to second order, complete Galilean invariant lattice BGK models in one dimension ($\theta = 3$) and two dimensions ($\theta = 1$) for the Navier-Stokes equation have been obtained.

Yeung, P.K., and Ye Zhou: *Numerical study of rotating turbulence with external forcing*. ICASE Report No. 98-39, (NASA/CR-1998-208702), August 31, 1998, 32 pages. Submitted to Physics of Fluids.

Direct numerical simulations at 256^3 resolution have been carried out to study the response of isotropic turbulence to the concurrent effects of solid-body rotation and numerical forcing at the large scales. Because

energy transfer to the smaller scales is weakened by rotation, energy input from forcing gradually builds up at the large scales, causing the overall kinetic energy to increase. At intermediate wavenumbers the energy spectrum undergoes a transition from a limited $k^{-5/3}$ inertial range to k^{-2} scaling recently predicted in the literature. Although the Reynolds stress tensor remains approximately isotropic and three-component, evidence for anisotropy and quasi- two-dimensionality in length scales and spectra in different velocity components and directions is strong. The small scales are found to deviate from local isotropy, primarily as a result of anisotropic transfer to the high wavenumbers. To understand the spectral dynamics of this flow we study the detailed behavior of nonlinear triadic interactions in wavenumber space. Spectral transfer in the velocity component parallel to the axis of rotation is qualitatively similar to that in non-rotating turbulence; however the perpendicular component is characterized by a greatly suppressed energy cascade at high wavenumbers and a local reverse transfer at the largest scales. The broader implications of this work are briefly addressed.

Hou, Thomas Y., Xiao-Hui Wu, Shiyi Chen, and Ye Zhou: *Effect of finite computational domain on turbulence scaling law in both physical and spectral spaces*. ICASE Report No. 98-40, (NASA/CR-1998-208710), September 11, 1998, 12 pages. Submitted to Physical Review E.

The well-known translation between the power law of energy spectrum and that of the correlation function or the second order structure function has been widely used in analyzing random data. Here, we show that the translation is valid only in proper scaling regimes. The regimes of valid translation are different for the correlation function and the structure function. Indeed, they do not overlap. Furthermore, in practice, the power laws exist only for a finite range of scales. We show that this finite range makes the translation inexact even in the proper scaling regime. The error depends on the scaling exponent. The current findings are applicable to data analysis in fluid turbulence and other stochastic systems.

Luo, Li-Shi: *A unified theory of non-ideal gas lattice Boltzmann models*. ICASE Report No. 98-41, (NASA/CR-1998-208711), September 11, 1998, 13 pages. To appear in Physical Review Letters.

A non-ideal gas lattice Boltzmann model is directly derived, in an *a priori* fashion, from the Enskog equation for dense gases. The model is rigorously obtained by a systematic procedure to discretize the Enskog equation (in the presence of an external force) in both phase space and time. The lattice Boltzmann model derived here is thermodynamically consistent and is free of the defects which exist in previous lattice Boltzmann models for non-ideal gases. The existing lattice Boltzmann models for non-ideal gases are analyzed and compared with the model derived here.

Montarnal, Philippe, and Chi-Wang Shu: *Real gas computations using an energy relaxation method and high-order WENO schemes*. ICASE Report No. 98-42, (NASA/CR-1998-208712), September 18, 1998, 26 pages. Submitted to the Journal of Computational Physics.

In this paper, we use a recently developed energy relaxation theory by Coquel and Perthame and high order weighted essentially non-oscillatory (WENO) schemes to simulate the Euler equations of real gas. The main idea is an energy decomposition into two parts: one part is associated with a simpler pressure law and the other part (the nonlinear deviation) is convected with the flow. A relaxation process is performed for each time step to ensure that the original pressure law is satisfied. The necessary characteristic decomposition for the high order WENO schemes is performed on the characteristic fields based on the first part. The

algorithm only calls for the original pressure law once per grid point per time step, without the need to compute its derivatives or any Riemann solvers. Both one and two dimensional numerical examples are shown to illustrate the effectiveness of this approach.

Wu, Jie-Zhi, Ye Zhou, and Meng Fan: *A note on kinetic energy, dissipation and enstrophy*. ICASE Report No. 98-43, (NASA/CR-1998-208713), September 18, 1998, 10 pages. Submitted to Physics of Fluids.

The dissipation rate of a Newtonian fluid with constant shear viscosity can be shown to include three constituents: dilatation, vorticity, and surface strain. The last one is found to make no contributions to the change of kinetic energy. These dissipation constituents are used to identify typical compact turbulent flow structures at high Reynolds numbers. The incompressible version of the simplified kinetic-energy equation is then cast to a novel form, which is free from the work rate done by surface stresses but in which the full dissipation reenters.

INTERIM REPORTS

Guerinoni, Fabio, Khaled Abdol-Hamid, and S. Paul Pao: *Parallel PAB3D: Experiences with a prototype in MPI*. ICASE Interim Report No. 31, (NASA/CR-1998-207636). April 3, 1998, 21 pages.

PAB3D is a three-dimensional Navier Stokes solver that has gained acceptance in the research and industrial communities. It takes as computational domain, a set disjoint blocks covering the physical domain. This is the first report on the implementation of PAB3D using the Message Passing Interface (MPI), a standard for parallel processing. We discuss briefly the characteristics of the code and define a prototype for testing. The principal data structure used for communication is derived from preprocessing "patching." We describe a simple interface (COMMSYS) for MPI communication, and some general techniques likely to be encountered when working on problems of this nature. Last, we identify levels of improvement from the current version and outline future work.

Brandt, Achi: *Barriers to achieving textbook multigrid efficiency (TME) in CFD*. ICASE Interim Report No. 32, (NASA/CR-1998-207647), May 19, 1998, 25 pages.

"Textbook multigrid efficiency" (TME) means solving a discrete PDE problem in a computational work which is only a small (less than 10) multiple of the operation count in the discretized system of equations itself. As a guide to attaining this optimal performance for general CFD problems, the table below lists every foreseen kind of computational difficulty for achieving that goal, together with the possible ways for resolving that difficulty, their current state of development, and references.

Included in the table are staggered and nonstaggered, conservative and nonconservative discretizations of viscous and inviscid, incompressible and compressible flows at various Mach numbers, as well as a simple (algebraic) turbulence model and comments on chemically reacting flows. The listing of associated computational barriers involves: non-alignment of streamlines or sonic characteristics with the grids; recirculating flows; stagnation points; discretization and relaxation on and near shocks and boundaries; far-field artificial boundary conditions; small-scale singularities (meaning important features, such as the complete airplane, which are not visible on some of the coarse grids); large grid aspect ratios; boundary layer resolution; and grid adaption.

ICASE COLLOQUIA
April 1, 1998 – September 30, 1998

Name/Affiliation/Title	Date
Homeier, Peter Vincent, University of Pennsylvania “Trustworthy Tools for Trustworthy Programs: A Mechanically Verified Verification Condition Generator for the Total Correctness of Procedures”	April 3
Dong, Jin-Song, CSIRO, Australia “Formal Object Modelling Techniques”	April 6
Toomer, Chris, Sowerby Research Centre, British Aerospace, Bristol, UK “Rapid Design Space Approximations for Aerodynamic Optimization”	May 12
Hovland, Paul, Argonne National Laboratory “Automatic Differentiation and Navier Stokes: In & Out”	May 14
Cousteix, Jean, ONERA, Toulouse, France “Laminar-turbulent Transition Activities at ONERA/DMAE”	May 26
Livingston, Mark, University of North Carolina at Chapel Hill “Tracking for Video-see-through Augmented Reality”	June 9
Choi, Kyung K., The University of Iowa - Center for Computer Aided Design “Design Velocity Field for Development of CAD-based Shape Design Process”	June 17
Iollo, Angelo, INRIA Sophia Antipolis, France “Approximation of Two-dimensional Flows by a Reduced Order Model”	June 23
Zannetti, Luca, DIAS Politecnico di Torino, Italy “Hamiltonian Modeling of Dispersion in Complex Topography”	June 24
Van Rosendale, John, National Science Foundation “Data and Visualization Corridors”	June 26
Barton, Russell, The Pennsylvania State University “A Metamodel Integration Strategy for System-level Design”	July 16
Renambot, Luc, INRIA/IRISA, France “Parallel Radiosity using Virtual Interfaces”	July 27

Name/Affiliation/Title	Date
Qian, Jian-Hua, National Center for Atmospheric Research "Toward High-resolution Climate Modeling"	August 7
Krause, Egon, Aerodynamisches Institut, RWTH, Germany "Flow Simulation for Hypersonic Vehicles"	August 13
Balsara, Dinshaw, NCSA, University of Illinois "The Design of Very High Order Accuracy Godunov Schemes for General Systems of Hyperbolic Equations (including MHD)"	August 14
Loncaric, Josip, ICASE "Spatial Structure of Distributed Flow Control"	August 17
Bokhari, Shahid, University of Engineering, Lahore, Pakistan "The Tera MTA and Unstructured Meshes"	August 21
Povitsky, Alexander, ICASE "Parallelization of Pipelined Algorithms in CFD"	August 21
Oliker, Leonid, RIACS, NASA Ames Research Center "PLUM: Parallel Load Balancing for Adaptive Unstructured Meshes"	August 24
Berkooz, Gal, BEAM Technologies, Inc. "PDESolve - An Object Oriented Environment for Development and Evolution of CAE Applications"	August 25
Drai, Remi, Universite de Nice-Sophia Antipolis, France "On the Linear Matrix Inequality Approach in Control"	August 27
Raveh, Daniella, Technion, Israel "Optimization of Flexible Aircraft Structure using Computational Fluid Dynamics"	August 31
Zhang, Xiaodong, The College of William & Mary "Exploiting Cache Locality at Runtime on Multiprocessor Systems"	September 10
Rauchwerger, Lawrence, Texas A&M University "Run-Time Parallelization: A Framework for Parallel Computation"	September 14
Zima, Hans, Institute for Software Technology and Parallel Systems, Austria "High Performance Fortran - Status and Future"	September 16

Name/Affiliation/Title	Date
Hart, William, Sandia National Laboratories “A Critical Review of Evolutionary Methods for Continuous Optimization”	September 23
Rossow, Cord -C., DLR, Germany “A Simple Flux Vector Splitting for Compressible Flows and its Extension to Incompressible Flows”	September 25

ICASE SUMMER ACTIVITIES

The summer program for 1998 included the following visitors:

<u>VISITOR and</u> <u>AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF</u> <u>VISIT</u>
Abarbanel, Saul <i>Applied & Numerical Math</i>	Tel Aviv University, Israel	8/10 – 9/04
Abdi, Frank <i>Applied & Numerical Math</i>	Alpha STAR Corporation	7/22 – 7/24
Baggag, Abdelkader <i>Computer Science</i>	Purdue University	6/15 – 9/11
Banks, H. Thomas <i>Applied & Numerical Math</i>	North Carolina State University	9/01 – 9/04
Batterman, Astrid <i>Applied & Numerical Math</i>	Universität Trier, Germany	5/04 – 5/30
Bokhari, Shahid <i>Computer Science</i>	Pakistan University of Engineering	7/13 – 9/04
Brandt, Achi <i>Applied & Numerical Math</i>	The Weizmann Institute of Science, Israel	7/06 – 7/10 8/05 – 8/14
Chamis, Christos <i>Applied & Numerical Math</i>	NASA Lewis Research Center	7/22 – 7/24
Chapman, Barbara <i>Computer Science</i>	University of Vienna, Austria	5/11 – 6/19
Chieng, Ching-Chang <i>Physical Sciences – Flow Control</i>	National Tsing Hua University, Taiwan	6/22 – 7/08
Cowan, F. Scott <i>Applied & Numerical Math</i>	Georgia Institute of Technology	6/15 – 8/21
Criminale, William <i>Physical Sciences – Fluid Mechanics</i>	University of Washington	9/04 – 9/25

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Darmofal, David <i>Applied & Numerical Math</i>	Massachusetts Institute of Technology	7/22 – 7/23 8/10 – 8/14
Diskin, Boris <i>Applied & Numerical Math</i>	The Weizmann Institute of Science, Israel	7/98 – 6/99
Edwards, Jack <i>Applied & Numerical Math</i>	North Carolina State University	7/15 – 8/10
Follett, William <i>Applied & Numerical Math</i>	Rocketdyne	7/20 – 7/24
Fu, Qiang <i>Physical Sciences Atmospheric Sciences</i>	University of Utah and Dalhousie University, Canada	5/26 – 7/24
Gantt, Christopher <i>Applied & Numerical Math</i>	Vanderbilt University	8/10 – 8/21
Gottlieb, David <i>Applied & Numerical Math</i>	Brown University	7/13 – 7/17 8/10 – 8/21 8/24 – 8/27
Grandhi, Ramana <i>Applied & Numerical Math</i>	Wright State University	7/21 – 7/24
Grosch, Chester <i>Physical Sciences – Fluid Mechanics</i>	Old Dominion University	5/18 – 8/18
Hayes, W. Scott <i>Computer Science</i>	The College of William & Mary	5/26 – 8/26
Higle, Julie <i>Applied & Numerical Math</i>	University of Arizona	7/22 – 7/24
Iollo, Angelo <i>Applied & Numerical Math</i>	INRIA, France	6/22 – 6/26

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Jackson, Tom <i>Physical Sciences – Fluid Mechanics</i>	University of Illinois - Urbana	8/31 – 9/04
Jen, Chun-Ping <i>Physical Sciences – Flow Control</i>	National Tsing Hua University, Taiwan	6/22 – 7/08
Keyes, David <i>Computer Science</i>	Old Dominion University	5/01 – 8/31
Lind, Richard <i>Applied & Numerical Math</i>	Veridian, Inc.	7/22 – 7/24
Llorente, Ignacio <i>Applied & Numerical Math</i>	Universidad Complutense, Spain	7/20 – 8/14
Longman, Richard <i>Physical Sciences – Flow Control</i>	Columbia University	5/26 – 5/29 6/22 – 8/22
Luetzgen, Gerald <i>Computer Science</i>	Universita at Passau, Germany	6/22 – 7/31
Mahadevan, Sankaran <i>Applied & Numerical Math</i>	Vanderbilt University	7/22 – 7/24
Manning, Valerie <i>Applied & Numerical Math</i>	Stanford University	8/04 – 8/21
Milder, Seth <i>Computer Science</i>	George Mason University	6/01 – 8/21
Muravyov, Alexander <i>Physical Sciences – Computational Structures</i>	University of British Columbia, Canada	6/15 – 8/31
Nordstrom, Jan <i>Applied & Numerical Math</i>	The Aeror auitcal Research Institute of Sweden	9/07 – 9/25
Ozturan, Can <i>Computer Science</i>	Bogazici University, Turkey	8/17 – 9/14

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Patten, James <i>Computer Science</i>	University of Virginia	5/27 – 8/18
Pothen, Alex <i>Computer Science</i>	Old Dominion University	5/11 – 8/15
Radespiel, Rolf <i>Applied & Numerical Math</i>	Institut für Entwurfsaerodynamik, (German Aerospace Center)	7/13 – 8/07
Rossow, -Ing. C.-C. <i>Applied & Numerical Math</i>	Institut für Entwurfsaerodynamik, (German Aerospace Center)	9/07 – 9/25
Ryaben'kii, Viktor <i>Applied & Numerical Math</i>	Russian Academy of Sciences	7/01 – 8/15
Sen, Suvrajeet <i>Applied & Numerical Math</i>	University of Arizona	7/22 – 7/24
Shu, Chi-Wang <i>Physical Sciences – Fluid Mechanics</i>	Brown University	7/06 – 7/13
Soemarwoto, Bambang <i>Applied & Numerical Math</i>	Institute of Technology Bandung, Indonesia	6/22 – 7/17
Sues, Robert H. <i>Applied & Numerical Math</i>	Applied Research Associate	7/22 – 7/24
Ta'asan, Shlomo <i>Applied & Numerical Math</i>	Carnegie Mellon University	5/04 – 5/29
Torczon, Virginia <i>Applied & Numerical Math</i>	The College of William & Mary	5/18 – 5/29 6/16 – 6/30
Trosset, Michael <i>Applied & Numerical Math</i>	University of Arizona	7/20 – 7/27
Tsynkov, Semyon <i>Applied & Numerical Math</i>	Tel Aviv University, Israel	6/29 – 9/25

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Turkel, Eli <i>Applied & Numerical Math</i>	Tel Aviv University, Israel	6/29 – 9/10
van Leer, Bram <i>Applied & Numerical Math</i>	The University of Michigan, Ann Arbor	6/22 – 8/07
Wallin, Stefan <i>Physical Sciences – Fluid Mechanics</i>	FFA - The Aeronautical Research Institute of Sweden	8/10 – 10/09
Xu, Kun <i>Applied & Numerical Math</i>	The Hong Kong University of Science and Technology	8/98 – 1/99
Yefet, Amir <i>Applied & Numerical Math</i>	Tel Aviv University, Israel	7/27 – 8/21
Zannetti, Luca <i>Applied & Numerical Math</i>	Politecnico di Torino, Italy	6/22 – 6/26
Zima, Hans <i>Computer Science</i>	Institute for Software Technology and Parallel Systems, Austria	8/17 – 9/18
Zubair, Mohammad <i>Computer Science</i>	Old Dominion University	6/01 – 8/31

OTHER ACTIVITIES

On April 22-24, 1998, ICASE/NASA LaRC/ARO/NSF co-sponsored a Workshop on Computational Aerosciences in the 21st Century at the Radisson Hotel in Hampton, VA. The objective of this workshop was to share new and emerging ideas that enable complex, multidisciplinary, numerical simulations in the aerospace sciences in time periods of minutes to hours instead of weeks or months. There were 96 attendees, and a formal proceedings will be published.

On August 3-4, 1998, ICASE and NASA LaRC co-sponsored an Aerostructure Workshop held at ICASE, NASA Langley Research Center. Methods for the steady or quasi-steady aerostructural interaction problem, i.e., aerodynamic performance, structural loads, and related issues such as roll effectiveness was the scope of this workshop. There were 21 attendees.

On August 18-20, 1998, ICASE and NASA Langley co-sponsored a Short Course on Wavelets and Multiresolution Analysis for Applications in the Aerosciences. This course was instructed by Professor Andrew J. Kurdila, University of Florida, Gainesville. The course presented a self-contained introduction to wavelets and multiresolution analysis, with a focus on engineering applications in the aerosciences. There were 50 attendees.

ICASE STAFF

I. ADMINISTRATIVE

Manuel D. Salas, Director, M.S., Aeronautics and Astronautics, Polytechnic Institute of Brooklyn, 1970.
Fluid Mechanics and Numerical Analysis.

Linda T. Johnson, Office and Financial Administrator

Etta M. Blair, Accounting Supervisor

Barbara A. Cardasis, Administrative Secretary

Rachel A. Lomas, Payroll and Accounting Clerk (Through May 1998)

Emily N. Todd, Conference Manager/Executive Assistant

Shannon K. Verstynen, Information Technologist

Gwendolyn W. Wesson, Contract Accounting Clerk

Shouben Zhou, Systems Manager (Beginning July 1998)

Leon M. Clancy, Senior System Manager (Through July 1998)

II. SCIENCE COUNCIL

Francine Berman, Professor, Department of Computer Science & Engineering, University of California-San Diego.

Joseph E. Flaherty, Amos Eaton Professor, Departments of Computer Science and Mathematical Sciences, Rensselaer Polytechnic Institute.

Geoffrey Fox, Director, Northeast Parallel Architectural Center, Syracuse University.

David Gottlieb, Professor, Division of Applied Mathematics, Brown University.

Forrester Johnson, Aerodynamics Research, Boeing Commercial Airplane Group.

Robert W. MacCormack, Professor, Department of Aeronautics and Astronautics, Stanford University.

Steven A. Orszag, Professor, Program in Applied and Computational Mathematics, Princeton University.

Stanley G. Rubin, Professor, Department of Aerospace Engineering and Engineering Mechanics, University of Cincinnati.

Manuel D. Salas, Director, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center.

III. RESEARCH FELLOWS

Dimitri Mavriplis - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Applied & Numerical Mathematics [Grid Techniques for Computational Fluid Dynamics]. (February 1997 to August 2001)

Piyush Mehrotra - Ph.D., Computer Science, University of Virginia, 1982. Computer Science [Programming Languages for Multiprocessor Systems]. (January 1991 to September 1999)

IV. CHIEF SCIENTIST

Geoffrey Lilley - Ph.D., Engineering, Imperial College, London, England, 1945. Fluid Mechanics. (April 1998 to December 1998)

V. SENIOR STAFF SCIENTISTS

Thomas W. Crockett - B.S., Mathematics, The College of William & Mary, 1977. Computer Science [System Software for Parallel Computing, Computer Graphics, and Scientific Visualization]. (February 1987 to August 2000)

Sharath S. Girimaji - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1990. Fluid Mechanics [Turbulence and Combustion]. (July 1993 to August 1999)

R. Michael Lewis - Ph.D., Mathematical Sciences, Rice University, 1989. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (May 1995 to August 2000)

Josip Lončarić - Ph.D., Applied Mathematics, Harvard University, 1985. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (March 1996 to August 1999)

Kwan-Liu Ma - Ph.D., Computer Science, University of Utah, 1993. Computer Science [Visualization]. (May 1998 to August 2001)

J. Ray Ristorcelli - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1991. Fluid Mechanics [Turbulence Modeling]. (December 1996 to June 1998)

David Sidilkover - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1989. Applied & Numerical Mathematics [Numerical Analysis and Algorithms]. (November 1994 to August 1999)

Ye Zhou - Ph.D., Physics, The College of William & Mary, 1987. Fluid Mechanics [Turbulence Modeling]. (October 1992 to September 1998)

VI. SCIENTIFIC STAFF

Brian G. Allan - Ph.D., Mechanical Engineering, University of California at Berkeley, 1996. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (February 1996 to August 1999)

Eyal Arian - Ph.D., Applied Mathematics, The Weizmann Institute of Science, Israel, 1995. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (October 1994 to August 1999)

Po-Shu Chen - Ph.D., Aerospace Engineering, University of Colorado-Boulder, 1997. Physical Sciences [Computational Structures]. (January 1998 to January 2001)

Stephen Gatterly - Ph.D., Computer Science, Carnegie Mellon University, 1995. Computer Science [Parallel Numerical Algorithms, including Partitioning and Mapping]. (September 1995 to August 1998)

Victoria L. Interrante - Ph.D., Computer Science, University of North Carolina at Chapel Hill, 1996. Computer Science [Scientific Visualization]. (March 1996 to July 1998)

Gerald Lüttgen - Ph.D., Computer Science, University of Passau, Germany, 1998. Computer Science [Formal Methods Research for Safety Critical Systems]. (October 1998 to August 2000)

Li-Shi Luo - Ph.D., Physics, Georgia Institute of Technology, 1993. Computer Science [Parallel Algorithms]. (November 1996 to October 1999)

Kwan-Liu Ma - Ph.D., Computer Science, University of Utah, 1993. Computer Science [Visualization]. (May 1993 to May 1998)

Alexander Povitsky - Ph.D., Mechanical Engineering, Moscow Institute of Steel and Alloys Technology (MISA), Russia, 1988. Computer Science [Parallelization and Formulation of Higher Order Schemes for Aeroacoustics Noise Propagation]. (October 1997 to August 2000)

VII. VISITING SCIENTISTS

Sang-Hyon Chu - Ph.D., Chemical Engineering, Seoul National University, 1998. Physical Sciences [Smart Materials and Flow Control]. (March 1998 to November 1998)

Boris Diskin - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1998. Applied & Numerical Mathematics [Convergence Acceleration]. (July 1998 to June 1999)

Remi Draï - M.S., Mathematics, Université de Nice-Sophia Antipolis, France, 1994. Teaching Assistant, Institut du Non linéaire de Nice and Department of Mathematics, Université de Nice-Sophia Antipolis, France. Physical Sciences [Modeling and Control Dealing with the Interaction of Fuel Slosh with Rigid Body Motion]. (March 1998 to October 1998)

Alexander Muravyov - Ph.D., Mechanical Engineering, University of British Columbia, Vancouver, 1997. Department of Mechanical Engineering, University of British Columbia, Vancouver. Fluid Mechanics [Computational Structural Acoustics]. (June 1998 to October 1998)

John Van Rosendale - Ph.D., Computer Science, University of Illinois, 1980. Program Director for New Technologies, Division of Advanced Scientific Computing, National Science Foundation. Computer Science [Parallel Computing]. (July 1994 to March 1999)

Hyeon-Ju Yoon - Ph.D., Computer Science, Korea Advanced Institute of Science and Technology, 1997. Postdoctoral Researcher, Department of Computer Science, Korea Advanced Institute of Science and Technology. Computer Science [Computer Graphics and Parallel and Distributed Processing]. (November 1997 to August 1998)

Kun Xu - Ph.D., Astrophysics, Columbia University, 1993. Assistant Professor, Department of Mathematics, The Hong Kong University of Science & Technology. Applied & Numerical Mathematics [Developing Gas-Kinetic Schemes] (August 1998 to January 1999)

VIII. SHORT TERM VISITING SCIENTISTS

Frank Abdi - Ph.D., Mechanical Engineering, University of Southern California, 1978. CEO, Alpha STAR Corporation, California. Optimization with Uncertainty. (July 1998)

Shahid H. Bokhari - Ph.D., Electrical and Computer Engineering, University of Massachusetts-Amherst, 1978. Professor, Department of Electrical Engineering, University of Engineering and Technology, Lahore, Pakistan. Computer Science [Mapping Computational Problems to New Parallel Architectures]. (July 1998 to September 1998)

Ching-Chang Chieng - Ph.D., Aerospace Engineering, Virginia Polytechnic Institute and State University, 1974. Professor, Department of Engineering and System Science, National Tsing Hua University, Taiwan. Physical Sciences [Controls]. (June 1998 to July 1998)

William O. Criminale - Ph.D., Aeronautics, The Johns Hopkins University, 1960. Professor, Department of Applied Mathematics, University of Washington. Fluid Mechanics. (September 1998)

Jack R. Edwards - Ph.D., Aerospace Engineering, North Carolina State University, 1993. Assistant Professor, Department of Aerospace Engineering, North Carolina State University. Applied & Numerical Mathematics [Convergence Acceleration]. (July 1998 to August 1998)

William W. Follett - M.S., Mechanical Engineering, Stanford University, 1987. Member of Technical Staff, Boeing, Rocketdyne Propulsion & Power. Applied & Numerical Mathematics. (July 1998)

Qiang Fu - Ph.D., Meteorology, University of Utah, 1991. Assistant Professor, Department of Oceanography, Dalhousie University, Nova Scotia. Atmospheric Science. (May 1998 to July 1998)

Ramana Grandhi - Ph.D., Engineering Science and Mechanics, Virginia Polytechnic Institute and State University, 1984. Distinguished Professor, Department of Mechanical and Materials, Wright State University. Optimization with Uncertainty. (July 1998)

Julia L. Higle - Ph.D., Industrial and Operations Engineering, University of Michigan, 1985. Associate Professor, Department of Systems and Industrial Engineering, University of Arizona. Optimization with Uncertainty. (July 1998)

Angelo Iollo - Ph.D., Aerospace Engineering, Politecnico di Torino, Italy, 1995. Assistant Professor of Fluid Mechanics, Department of Aerospace Engineering, Politecnico di Torino, Italy. Applied & Numerical Mathematics. (June 1998)

Thomas L. Jackson - Ph.D., Mathematics, Rensselaer Polytechnic Institute, 1985. Senior Research Scientist, Center for Simulation of Advanced Rockets, University of Illinois. Physical Sciences [Fluid Mechanics]. (September 1998)

Ignacio M. Llorente - Ph.D., Computer Science, Complutense University of Madrid, Spain, 1995. Associate Professor, Department of Computer Architecture, Complutense University of Madrid, Spain. Computer Science. (July 1998 to August 1998)

Gerald Lüttgen - Ph.D., Computer Science, University of Passau, Germany, 1998. Staff Scientist, Department of Mathematics and Computer Science, University of Passau, Germany. Computer Science [Safety Critical Systems]. (June 1998 to September 1998)

Sankaran Mahadevan - Ph.D., Civil Engineering, Georgia Institute of Technology, 1988. Associate Professor, Department of Civil and Environmental Engineering, Vanderbilt University. Optimization with Uncertainty. (July 1998)

Can Özturan - Ph.D., Computer Science, Rensselaer Polytechnic Institute, 1995. Assistant Professor, Department of Computer Engineering, Bogazici University, Istanbul, Turkey. Computer Science. (August 1998 to September 1998)

Rolf Radespiel - Ph.D., Engineering, University of Braunschweig, Germany, 1986. Branch Head, DLR, Institute of Design Aerodynamics, Germany. Applied & Numerical Mathematics. (July 1998 to August 1998)

Cord-Christian Rossow - Ph.D., Aerospace Engineering, Technical University of Braunschweig, Germany, 1988. Branch Head, Dr.-Ing, DLR, Institute of Design Aerodynamics, Germany. Applied & Numerical Mathematics. (September 1998)

Suvrajeet Sen - Ph.D., Operations Research, Virginia Polytechnic Institute and State University, 1982. Professor, Department of Systems and Industrial Engineering, University of Arizona. Optimization with Uncertainty. (July 1998)

Bambang I. Soemarwoto - Ph.D., Aeronautics, Delft University of Technology, The Netherlands, 1996. Lecturer in Aerodynamics, Department of Aeronautics & Astronautics, Institute of Technology, Bandung, Indonesia. Applied & Numerical Mathematics. (June 1998 to July 1998)

Robert H. Sues - Ph.D., Structural Engineering, University of Illinois, 1983. Principal Engineer, Applied Research Associates, Raleigh, NC. Optimization with Uncertainty. (July 1998)

Shlomo Ta'asan - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1985. Professor, Department of Mathematics, Carnegie Mellon University. Applied & Numerical Mathematics. (May 1998)

Michael W. Trosset - Ph.D., Statistics, University of California-Berkeley, 1983. Visiting Associate Professor, Department of Mathematics, University of Arizona. Applied & Numerical Mathematics. (July 1998)

Eli Turkel - Ph.D., Applied Mathematics, New York University, 1970. Associate Professor, Department of Applied Mathematics, Tel Aviv University, Israel. Applied & Numerical Mathematics. (June 1998 to September 1998)

O. Stefan Wallin - Ph.D., Mechanical Engineering, Technical University of Lulea, Sweden. Research Engineer, Department of Computational Aerodynamics, The Aeronautical Research Institute of Sweden. Fluid Mechanics. (August 1998 to October 1998)

Luca Zannetti - Ph.D., Aerospace Engineering, Politecnico di Torino, 1971. Professor, Dipartimento di Ingegneria Aeronautica, Politecnico di Torino, Italy. Applied & Numerical Mathematics. (June 1998)

IX. ASSOCIATE RESEARCH FELLOW

David E. Keyes - Ph.D., Applied Mathematics, Harvard University, 1984. Computer Science [Parallel Numerical Algorithms]

X. CONSULTANTS

Saul Abarbanel - Ph.D., Theoretical Aerodynamics, Massachusetts Institute of Technology, 1959. Professor, Department of Applied Mathematics, Tel Aviv University, Israel. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamics and Aeroacoustic Computations]

Ponnampalam Balakumar - Ph.D., Aeronautics and Astronautics, Massachusetts Institute of Technology, 1986. Associate Professor, Department of Aerospace Engineering, Old Dominion University. Fluid Mechanics [Stability and Transition]

H. Thomas Banks - Ph.D., Applied Mathematics, Purdue University, 1967. Professor, Department of Mathematics, Center for Research in Scientific Computations, North Carolina State University. Applied & Numerical Mathematics [Control Theory]

Richard W. Barnwell - Ph.D., Engineering Mechanics, Virginia Polytechnic Institute and State University, 1968. Professor, Department of Aerospace and Ocean Engineering, Engineering Science and Mechanics. Virginia Polytechnic Institute and State University. Fluid Mechanics [Turbulence Modeling]

Alvin Bayliss - Ph.D., Mathematics, New York University, 1975. Associate Professor, Technological Institute, Northwestern University. Fluid Mechanics [Numerical Solution of the Equations of Fluid Flow and Acoustics]

Oktay Baysal - Ph.D., Mechanical Engineering, Louisiana State University, 1982. Eminent Scholar and Professor, Department of Aerospace Engineering, Old Dominion University. Applied & Numerical Mathematics

Achi Brandt - Ph.D., Mathematics, Weizmann Institute of Science, 1965. Professor, Department of Applied Mathematics, Weizmann Institute of Science, Israel. Applied & Numerical Mathematics [Convergence Acceleration]

Barbara M. Chapman - M.S., Mathematics, University of Canterbury, Christchurch, New Zealand, 1985. Director, European Institute for Parallel Computing, University of Vienna. Computer Science [Parallel Language Extensions and Optimizations for Parallel Compilers]

Ayodeji O. Demuren - Ph.D., Mechanical Engineering, Imperial College London, United Kingdom, 1979. Associate Professor, Department of Mechanical Engineering and Mechanics, Old Dominion University. Fluid Mechanics [Numerical Modeling of Turbulent Flows]

Geoffrey Fox - Ph.D., Physics, Cambridge University, 1967. Professor, Department of Computer Science, Syracuse University. Computer Science [Networking]

David Gottlieb - Ph.D., Numerical Analysis, Tel-Aviv University, Israel, 1972. Ford Foundation Professor and Chair, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Boundary Conditions for Hyperbolic Systems]

Chester E. Grosch - Ph.D., Physics and Fluid Dynamics, Stevens Institute of Technology, 1967. Professor, Department of Computer Science and Slover Professor, Department of Oceanography, Old Dominion University. Fluid Mechanics [Turbulence and Acoustics]

Jan S. Hesthaven - Ph.D., Applied Mathematics/Numerical Analysis, Technical University of Denmark, 1995. Visiting Assistant Professor, Division of Applied Mathematics, Brown University. Physical Sciences [Computational Electromagnetics]

Fang Q. Hu - Ph.D., Applied Mathematics, Florida State University, 1990. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Aeroacoustics]

Frank Kozusko - Ph.D., Computational and Applied Mathematics, Old Dominion University, 1995. Assistant Professor, Department of Mathematics, Hampton University. Fluid Mechanics [Airfoil Design]

David G. Lasseigne - Ph.D., Applied Mathematics, Northwestern University, 1985. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Asymptotic and Numerical Methods for Computational Fluid Dynamics]

Richard W. Longman - Ph.D., Aerospace Engineering, University of California-San Diego, 1969. Professor, Department of Mechanical Engineering, Columbia University. Physical Science [System/Disturbance Identification for Flow Control]

Kurt Maly - Ph.D., Computer Science, Courant Institute, New York University, 1973. Kaufman Professor and Chair, Department of Computer Science, Old Dominion University. Computer Science [High Performance Communication]

James E. Martin - Ph.D., Applied Mathematics, Brown University, 1991. Assistant Professor, Department of Mathematics, Christopher Newport University. Fluid Mechanics [Turbulence and Computation]

Jan Nordstrom - Ph.D., Numerical Analysis, Uppsala University, Sweden, 1993. Senior Scientist, The Aeronautical Research Institute of Sweden. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamic and Aeroacoustic Computations]

Alex Pothén - Ph.D., Applied Mathematics, Cornell University, 1984. Professor, Department of Computer Science, Old Dominion University. Computer Science [Parallel Numerical Algorithms]

Robert Rubinstein - Ph.D., Mathematics, Massachusetts Institute of Technology, 1972. Fluid Mechanics [Drag Reduction and Shock Wave Propagation in a Gas Far-From Equilibrium]

Viktor Ryaben'kii - Ph.D., Stability of Difference Equations, Moscow State University, 1953. Leading Research Scientist, Keldysh Institute for Applied Mathematics, Russian Academy of Sciences and Full Professor, Department of Control and Applied Mathematics, Moscow Institute of Physics and Technology. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamic and Aeroacoustic Computations]

Chi-Wang Shu - Ph.D., Mathematics, University of California-Los Angeles, 1986. Associate Professor, Division of Applied Mathematics, Brown University. Fluid Mechanics [Computational Aeroacoustics]

Ralph C. Smith - Ph.D., Mathematics, Montana State University, 1990. Assistant Professor, Department of Mathematics, Iowa State University. Applied & Numerical Mathematics [Optimal Control Techniques for Structural Acoustics Problems]

Siva Thangam - Ph.D., Mechanical Engineering, Rutgers University, 1980. Professor, Department of Mechanical Engineering, Stevens Institute of Technology. Fluid Mechanics [Turbulence Modeling and Simulation]

Christine A. Toomer - Ph.D., Mathematics, Imperial College of Science, Technology & Medicine, London University, 1990. Research Scientist, Aerodynamics & Vulnerability Department, Sowerby Research Centre, British Aerospace, United Kingdom. Applied & Numerical Mathematics [Methods for Aerodynamic Design and Optimization]

Virginia Torczon - Ph.D., Mathematical Sciences, Rice University, 1989. Assistant Professor, Department of Computer Science, The College of William & Mary. Computer Science [Parallel Algorithms for Optimization including Multidisciplinary Optimization]

Michael W. Trosset - Ph.D., Statistics, University of California-Berkeley, 1983. Department of Mathematics, The College of William & Mary. Applied & Numerical Mathematics [Multidisciplinary Optimization]

Semyon V. Tsynkov - Ph.D., Computational Mathematics, Keldysh Institute for Applied Mathematics, Russian Academy of Sciences, 1991. Senior Lecturer, Department of Applied Mathematics, Tel Aviv University, Israel. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamic and Aeroacoustic Computations]

George M. Vahala - Ph.D., Physics, University of Iowa, 1972. Professor, Department of Physics, The College of William & Mary. Fluid Mechanics [Group Renormalization Methods for Turbulence Approximation]

Bram van Leer - Ph.D., Theoretical Astrophysics, Leiden State University, The Netherlands, 1970. Professor, Department of Aerospace Engineering, University of Michigan. Applied & Numerical Mathematics [Convergence Acceleration]

Robert G. Voigt - Ph.D., Mathematics, University of Maryland, 1969. Professor, Computational Science Program, The College of William & Mary. Computer Science [High Performance Computing]

Hans Zima - Ph.D., Mathematics, University of Vienna, Austria, 1964. Professor, Institute for Software Technology and Parallel Systems, University of Vienna, Austria. Computer Science [Compiler Development for Parallel and Distributed Multiprocessors]

Mohammad Zubair - Ph.D., Computer Science, Indian Institute of Technology, Delhi, India, 1987. Professor, Department of Computer Science, Old Dominion University. Computer Science [Performance of Unstructured Flow-Solvers on Multi Processor Machines]

XI. GRADUATE STUDENTS

Ahmed H. Al-Thencyan - Department of Computer Science, Old Dominion University. (February 1998 to August 25, 1998)

Abdelkader Baggag - Department of Computer Science, Purdue University. (September 1995 to Present)

Astrid Batterman - Department of Mathematics, Universitaet Trier, Germany. (May 1998)

Ya-Chin Chen - Department of Electrical and Electronic Engineering, Imperial College of Science, Technology, and Medicine, London, United Kingdom. (October 1997 to August 1998)

F. Scott Cowan - Systems Realization Laboratory, G.W.W. School of Mechanical Engineering, Georgia Institute of Technology. (June 1998 to August 1998)

David C. Cronk - Department of Computer Science, The College of William & Mary. (August 1993 to October 1998)

Elizabeth D. Dolan - Department of Computer Science, The College of William & Mary. (August 1998 to Present)

Christopher W. Gantt - Department of Civil and Environmental Engineering, Vanderbilt University. (August 1998)

W. Scott Hayes - Department of Computer Science, The College of William & Mary. (May 1998 - September 1998)

David A. Hysom - Department of Computer Science, Old Dominion University. (October 1997 to Present)

Chung-Ping Jen - Department of Engineering and System Science, National Tsing Hua University, Hsin Chu, Taiwan. (June 1998 to July 1998)

Nilan Karunaratne - Department of Computer Science, Old Dominion University. (August 1995 to Present)

Dinesh Kaushik - Department of Computer Science, Old Dominion University. (May 1997 to Present)

Gary Kumfert - Department of Computer Science, Old Dominion University. (January 1997 to Present)

Georgia Liu - Department of Computer Science, Old Dominion University. (February 1997 to Present)

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY(Leave blank)	2. REPORT DATE December 1998	3. REPORT TYPE AND DATES COVERED Contractor Report		
4. TITLE AND SUBTITLE Semiannual Report. April 1, 1998 through September 30, 1998		5. FUNDING NUMBERS C NAS1-97046 WU 505-90-52-01		
6. AUTHOR(S)				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Computer Applications in Science and Engineering Mail Stop 403, NASA Langley Research Center Hampton, VA 23681-2199		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Langley Research Center Hampton, VA 23681-2199		10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA/CR-1998-208959		
11. SUPPLEMENTARY NOTES Langley Technical Monitor: Dennis M. Bushnell Final Report				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified-Unlimited Subject Category 59 Distribution: Nonstandard Availability: NASA-CASI (301)621-0390		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) This report summarizes research conducted at the Institute for Computer Applications in Science and Engineering in applied mathematics, fluid mechanics, and computer science during the period April 1, 1998 through September 30, 1998.				
14. SUBJECT TERMS applied mathematics; multidisciplinary design optimization; fluid mechanics; turbulence; flow control; acoustics; computer science; system software; systems engineering; parallel algorithms			15. NUMBER OF PAGES 78	
			16. PRICE CODE A05	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	

Valerie Manning - Department of Aeronautics and Astronautics, Stanford University. (August 1998)

Seth D. Milder - Department of Physics and Astronomy, George Mason University. (September 1997 to Present)

James Patten - Department of Computer Science, University of Virginia. (May 1998 to August 1998)

Deborah F. Pilkey - Department of Engineering Sciences and Mechanics, Virginia Polytechnic Institute and State University. (October 1995 to May 1998)

Kevin Roe - Department of Computer Science, The College of William & Mary. (May 1995 to Present)

Ming-Yun Shih - Department of Computer Science, The College of William & Mary. (May 1997 to Present)

Diann P. Smith - Department of Computer Science, The College of William & Mary. (August 1997 to July 1998)

Amir Yefet - School of Mathematical Sciences, Tel Aviv University, Israel. (July 1998 to August 1998)